




FDR®

TOTAL STORAGE
MANAGEMENT SYSTEM

V5.3



INNOVATION®
DATA PROCESSING

CORPORATE HEADQUARTERS: 275 Paterson Ave., Little Falls, NJ 07424 • (973) 890-7300 • Fax: (973) 890-7147
E-mail: support@fdrinnovation.com • sales@fdrinnovation.com • <http://www.innovationdp.fdr.com>

EUROPEAN OFFICES:	FRANCE	GERMANY	NETHERLANDS	UNITED KINGDOM	NORDIC COUNTRIES
	01-49-69-94-02	089-489-0210	036-534-1660	0208-905-1266	+31-36-534-1660

FAST DUMP RESTORE (FDR) USER MANUAL

VERSION V5.3

VOLUME 1

PURPOSE OF THIS MANUAL

This manual is in two binders and provides the information you need to use and understand the FDR DASD MANAGEMENT SYSTEM, which is composed of these functional components (some of which are optional and separately priced):

VOLUME 1

- **FDR** (FAST DUMP RESTORE) for full volume backup, restore and disk-to-disk copy
- **DSF** (DATA SET FUNCTIONS) for data set backup and restore, including **FDRCOPY** for disk-to-disk copy
- **SAR** (STAND ALONE RESTORE) for backup and restore without an operating system
- **FDR INSTANTBACKUP** for instant backups and reorganizations when used with special hardware features
- **FDRREORG** for automating reorganization of VSAM KSDS, IAM, and PDS data sets
- **FASTCPK** (COMPAKTOR) - for DASD volume reorganization.
- **ABR** (AUTOMATIC BACKUP & RECOVERY) for automation of backups, restores, and space management
- **FDRCLONE** for quickly "cloning" volumes and/or data sets on a test or disaster system. It includes **FDRDRP**, a utility which can reduce full-volume recovery time up to 80%.

VOLUME 2

- **FDRAPPL** for backup of data belonging to a given application from whatever volume it resides on.
- **FDREPORT** for high-performance, flexible, and powerful reporting on DASD data and related information

Volume 2 also contains

other FDR reports,
FDR utilities **FDRTCOPY** and **FDRTSEL**
SMS considerations
other considerations
installation instructions
and FDR messages.

INNOVATION DATA PROCESSING

FDR, DSF, COMPAKTOR, ABR, AUTOMATIC BACKUP & RECOVERY, FDRCLONE, FDRREORG and FDREPORT are trademarks of INNOVATION DATA PROCESSING and are registered with the US patent office.

FDR[®] is a proprietary program products of
INNOVATION DATA PROCESSING
Innovation Plaza
275 Paterson Avenue
Little Falls, New Jersey 07424-1658

and are available for license exclusively from INNOVATION DATA PROCESSING, INC.

© Copyright 1972, 1973, 1976, 1981, 1982, 1983, 1984, 1986, 1988, 1989, 1990, 1991, 1992, 1993, 1998, 1999, 2000

© INNOVATION DATA PROCESSING, INC.

Reproduction of this User Manual is prohibited except by licensed users for their internal use.

A copy of this manual is provided on CD-ROM in IBM BookManager format for online use and in Adobe PDF format for local printing.

US CONTACTS

Voice: 973-890-7300 Fax: 973-890-7147
Home page: www.innovationdp.fdr.com

E-mail: Sales: sales@fdrinnovation.com
Technical Support: support@fdrinnovation.com

See next page for Innovation Offices in Europe.

FAST DUMP RESTORE (FDR) USER MANUAL

CONTINUED

EUROPEAN OFFICES:

France

Innovation Data Processing S.A.R.L.
191 Avenue Aristide Briand
94230 CACHAN
Tel: 01-49-69-94-02 Fax: 01-49-69-90-98
E-mail: frsupport@fdrinnovation.com
frsales@fdrinnovation.com

Germany

Innovation Data Processing
International Ltd.
Orleansstrasse 4a
D-81669 München
Tel: 089-489 0210 Fax: 089-489 1355
E-mail: desupport@fdrinnovation.com
desales@fdrinnovation.com

Netherlands

Innovation Data Processing
Brouwerstraat 8
1315 BP Almere
Tel: 036-534 1660 Fax: 036-533 7308
E-mail: nlsupport@fdrinnovation.com
nlsales@fdrinnovation.com

United Kingdom

Innovation Data Processing Ltd.
Clarendon House
125 Shenley Road
Borehamwood, Herts, WD6 1AG
Tel: 0208-905 1266 Fax: 0208-905 1428
E-mail: uksupport@fdrinnovation.com
uksales@fdrinnovation.com

CONTINUED . . .

TABLE OF CONTENTS

Page No.

01.0 ORGANIZATION OF THIS MANUAL

01.01	ORGANIZATION OF THIS MANUAL	xv
01.02	CONTROL STATEMENT FORMAT	xvi
01.03	YEAR 2000 COMPLIANCE	xix
01.04	SUMMARY OF MODIFICATIONS FOR VERSION 5.3	xxi
01.04	SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 50	xxi
01.04	SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 40	xxii
01.04	SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 30	xxv
01.04	SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 22	xxvii
01.04	SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 20	xxviii

02.0 FDR DASD MANAGEMENT SYSTEM

02.01	FDR DASD MANAGEMENT SYSTEM	02-1
02.02	FDR AND FDRDSF INTRODUCTION	02-3
02.03	COMPAKTOR INTRODUCTION	02-7
02.04	FDRABR INTRODUCTION	02-8
02.05	FDREPORT INTRODUCTION	02-12
02.06	FDRREORG INTRODUCTION	02-13
02.10	FDR FUNCTIONS AND FEATURES	02-14

10.0 FDR OVERVIEW

10.01	FDR OVERVIEW	10-1
10.02	FDR PROCESSING OPTIONS AND REQUIREMENTS	10-3
10.03	FDR JCL REQUIREMENTS	10-9
10.04	FDR DUMP STATEMENT	10-14
10.05	FDR RESTORE STATEMENT	10-19
10.06	FDR COPY STATEMENT	10-23
10.10	FDR DUMP EXAMPLES	10-28
10.11	FDR RESTORE EXAMPLES	10-32
10.12	FDR DISK COPY EXAMPLES	10-35
10.13	FDR INSTANTBACKUP EXAMPLES	10-37
10.25	FDRQUERY - FDR/ABR STATISTICS QUERY	10-39
10.26	FDRQUERY JOB CONTROL	10-40
10.27	FDRQUERY REPORT STATEMENT	10-41
10.28	FDRQUERY EXAMPLES	10-43
10.29	FDRQUERY SUMMARY ILLUSTRATION	10-44

15.0 SAR TECHNICAL DESCRIPTION

15.01	SAR TECHNICAL DESCRIPTION	15-1
15.02	SAR OPERATION REQUIREMENTS	15-6
15.03	LOADING SAR	15-8
15.04	SAR CONSOLE MESSAGES	15-11
15.20	STAND-ALONE LOADER UTILITY	15-16
15.21	FDRSARLR JOB CONTROL REQUIREMENTS	15-17
15.22	FDRSARLR CONTROL STATEMENTS	15-19
15.23	FDRSARLR EXAMPLES	15-20
15.24	FDRSARLR ISPF INTERFACE	15-22
15.25	SAR HINTS AND RECOMMENDATIONS	15-23

TABLE OF CONTENTS

Page No.

20.0	DSF TECHNICAL SUMMARY 1	
20.01	DSF TECHNICAL SUMMARY	20-1
20.02	DSF PROCESSING OPTIONS, PERFORMANCE AND REQUIREMENTS ...	20-7
20.03	DSF DUMP/PRINT JOB CONTROL REQUIREMENTS	20-12
20.04	DSF DUMP STATEMENT	20-16
20.05	DSF PRINT STATEMENT	20-22
20.06	DSF SELECT STATEMENT -- FOR DUMP OR PRINT	20-23
20.07	DSF RESTORE JOB CONTROL REQUIREMENTS	20-26
20.08	DSF RESTORE STATEMENT	20-28
20.09	DSF SELECT STATEMENT -- FOR RESTORES	20-35
20.10	DSF DUMP EXAMPLES	20-44
20.11	DSF PRINT EXAMPLES	20-48
20.12	DSF NON-SMS RESTORE EXAMPLES	20-49
20.13	DSF SMS RESTORE EXAMPLES	20-53
20.14	DSF UNLIKE DEVICE EXAMPLES	20-56
20.15	FDR INSTANTBACKUP EXAMPLES	20-58
21.0	FDRCOPY OVERVIEW	
21.01	FDRCOPY OVERVIEW	21-1
21.02	FDRCOPY PROCESSING OPTIONS AND REQUIREMENTS	21-9
21.03	FDRCOPY JOB CONTROL REQUIREMENTS	21-13
21.04	FDRCOPY COPY/MOVE STATEMENT	21-15
21.05	FDRCOPY SELECT/EXCLUDE STATEMENT FOR COPY/MOVE	21-21
21.10	FDRCOPY DATA SET COPY EXAMPLES	21-31
21.11	FDRCOPY DATA SET MOVE EXAMPLES	21-32
21.12	FDRCOPY SMS EXAMPLES	21-34
21.13	FDRCOPY UNLIKE DEVICE EXAMPLES	21-36
21.14	FDR INSTANTBACKUP EXAMPLES	21-37
21.20	FDRCOPY ISPF INTERFACE	21-38
25.0	FDR INSTANTBACKUP FOR EMC TIMEFINDER™	
25.01	FDR INSTANTBACKUP FOR EMC TIMEFINDER™	25-1
25.02	TIMEFINDER OPERATION	25-2
25.03	THE FDR INSTANTBACKUP SOLUTION	25-3
25.04	SETTING UP BCVS	25-4
25.05	FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY	25-7
25.06	FDR INSTANTBACKUP FOR ABR	25-10
25.07	EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK	25-13
25.08	EXAMPLES FOR ABR	25-17

TABLE OF CONTENTS

Page No.

26.0	FDR INSTANTBACKUP FOR SNAPSHOT	
26.01	FDR INSTANTBACKUP FOR SNAPSHOT	26-1
26.02	SNAPSHOT OPERATION	26-2
26.03	THE FDR INSTANTBACKUP SOLUTION	26-4
26.04	SETTING UP SNAPSHOT	26-5
26.05	FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY	26-7
26.06	FDR INSTANTBACKUP FOR ABR	26-9
26.07	EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK	26-12
26.08	EXAMPLES FOR ABR	26-16
27.0	FDR INSTANTBACKUP FOR HITACHI DATA SYSTEMS SHADOWIMAGE	
27.01	FDR INSTANTBACKUP FOR HITACHI DATA SYSTEMS SHADOWIMAGE ..	27-1
27.02	SHADOWIMAGE OPERATION	27-2
27.03	THE FDR INSTANTBACKUP SOLUTION	27-4
27.04	SETTING UP SHADOWIMAGE VOLUMES	27-5
27.05	FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY	27-8
27.06	FDR INSTANTBACKUP FOR ABR	27-11
27.07	EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK	27-14
27.08	EXAMPLES FOR ABR	27-17
28.0	FDR INSTANTBACKUP FOR AMDAHL TDMF⁴	
28.01	FDR INSTANTBACKUP FOR AMDAHL TDMF ⁴	28-1
28.02	TDMF OPERATION	28-2
28.03	THE FDR INSTANTBACKUP SOLUTION	28-3
28.04	SETTING UP TDMF	28-4
28.05	FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY	28-5
28.06	EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK	28-8
29.0	FDR INSTANTBACKUP FOR IBM FLASHCOPY	
29.01	FDR INSTANTBACKUP FOR IBM FLASHCOPY	29-1
29.02	FLASHCOPY OPERATION	29-2
29.03	THE FDR INSTANTBACKUP SOLUTION	29-4
29.04	SETTING UP FLASHCOPY	29-5
29.05	FDR INSTANTBACKUP FOR ABR	29-7
29.06	FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY	29-10
29.07	EXAMPLES FOR FDR/FDRDSF/FDRCOPY	29-13
29.08	EXAMPLES FOR ABR	29-16

TABLE OF CONTENTS

Page No.

30.0	FDRREORG OVERVIEW	
30.01	FDRREORG OVERVIEW	30-1
30.02	FDRREORG JOB CONTROL REQUIREMENTS AND DEFAULTS	30-11
30.03	FDRREORG REORG AND SIMULATE STATEMENTS	30-12
30.04	FDRREORG EXCLUDE STATEMENT	30-28
30.05	FDRREORG SELECT STATEMENT	30-32
30.06	FDRREORG RECOVER STATEMENT	30-46
30.07	FDRREORG OPERATOR COMMANDS	30-54
30.10	FDRREORG SIMULATE EXAMPLES	30-55
30.11	FDRREORG EXAMPLES	30-58
30.12	FDRREORG RECOVER EXAMPLES	30-67
31.0	FDRREORG OPTIONS AND DEFAULTS	
31.01	FDRREORG OPTIONS AND DEFAULTS	31-1
31.01	FDRREOZO JCL REQUIREMENTS	31-2
31.03	FDRREOZO STATEMENTS	31-3
31.04	FDRREOZO NOREORG STATEMENTS	31-4
31.05	FDRREOZO ZAP STATEMENT	31-5
31.06	FDRREORG JCL EXAMPLES	31-16
32.0	FDRCOPY REORG STATEMENT	
32.01	FDRCOPY REORG STATEMENT	32-1
32.02	FDRCOPY SELECT/EXCLUDE STATEMENT FOR REORG	32-4
32.03	FDRCOPY REORG EXAMPLES	32-8
33.0	FDRREORG ISPF INTERFACE	
33.01	FDRREORG ISPF INTERFACE	33-1

TABLE OF CONTENTS

Page No.

40.0	COMPAKTOR INTRODUCTION	
40.01	COMPAKTOR INTRODUCTION	40-1
40.02	COMPAKTOR BACKGROUND	40-2
40.03	COMPAKTOR OVERVIEW	40-3
40.04	COMPAKTOR SPECIAL FEATURES	40-5
40.05	COMPAKTOR JCL REQUIREMENTS	40-8
40.06	COMPAKTOR STATEMENTS	40-10
40.07	MAP STATEMENT	40-11
40.08	COMPAKT AND SIMULATE STATEMENTS	40-14
40.09	SELECT STATEMENT	40-28
40.10	VTOC SELECT STATEMENT	40-32
40.11	SEQUENCE AND ENDSEQ STATEMENT	40-34
40.12	COMPAKTOR DATA SET TYPES	40-35
40.13	COMPAKTOR DATA SET POSITIONING	40-37
40.14	COMPAKTOR REORGANIZATION	40-38
40.15	COMPAKTOR OUTPUT	40-40
40.16	COMPAKTOR TECHNICAL DESCRIPTION	40-41
40.17	COMPAKTOR VTOC ERRORS	40-44
40.18	COMPAKTOR ERROR RECOVERY	40-45
40.19	COMPAKTOR PROBLEM DETERMINATION	40-46
40.20	RECOVERING FROM COMPAKTOR ERRORS	40-47
40.21	COMPAKTING ACTIVE VOLUMES	40-50
40.22	COMPAKTOR REPORTS	40-52
40.23	COMPAKTOR EXAMPLES	40-57
40.24	COMPAKTOR ISPF DIALOGS	40-64
40.25	CPK HINTS AND RECOMMENDATIONS	40-66
40.26	BENEFITS OF FASTCPK ON VIRTUAL DISK ARRAYS	40-71

TABLE OF CONTENTS

Page No.

50.0	FDRABR VOLUME BACKUPS	
50.01	FDRABR VOLUME BACKUPS	50-1
50.02	DATA SET SELECTION	50-10
50.03	PROCESSING OPTIONS AND REQUIREMENTS	50-14
50.04	VOLUME BACKUP JOB CONTROL REQUIREMENTS	50-22
50.05	VOLUME BACKUP DUMP/SIM/SNAP/SPLIT/PSPLIT STATEMENTS	50-27
50.06	VOLUME BACKUP SELECT, EXCLUDE AND MOUNT STATEMENTS	50-38
50.07	RESTORE JOB CONTROL REQUIREMENTS	50-44
50.08	DATA SET RESTORE STATEMENT	50-47
50.09	SELECT STATEMENT FOR DATA SET RESTORE	50-55
50.10	FULL VOLUME RESTORE STATEMENT	50-67
50.11	SELECT STATEMENT FOR FULL VOLUME RESTORE	50-72
50.20	VOLUME BACKUP EXAMPLES	50-74
50.21	DATA SET RESTORE EXAMPLES	50-78
50.22	SMS DATA SET RESTORE EXAMPLES	50-81
50.23	FULL VOLUME RESTORE EXAMPLES	50-82
50.24	FDR INSTANTBACKUP EXAMPLES	50-83
50.30	VOLUME BACKUP ISPF SUPPORT	50-91
50.40	VTOC MAINTENANCE UTILITY (FDRABRM)	50-95
50.41	FDRABRM JCL REQUIREMENTS	50-98
50.42	FDRABRM ABRINIT STATEMENT	50-99
50.43	FDRABRM MAINT STATEMENT	50-104
50.44	FDRABRM REMODEL STATEMENT	50-105
50.45	VTOC UTILITY JCL EXAMPLES	50-107
50.46	FDRABRM ISPF INTERFACE	50-109
50.50	CATALOG MAINTENANCE UTILITY (FDRABRCM)	50-111
50.51	FDRABRCM JCL REQUIREMENTS	50-112
50.52	FDRABRCM PURGE SCRATCH STATEMENT	50-113
50.53	FDRABRCM PURGE BACKUP STATEMENT	50-116
50.54	FDRABRCM EXAMPLES	50-119
50.70	FDRCLONE	50-121
50.71	FDRCLONE OPERATION	50-125
50.72	FDRCLONE JOB CONTROL REQUIREMENTS	50-131
50.73	FDRCLONE CLONE STATEMENT	50-132
50.74	FDRCLONE SELECT AND EXCLUDE STATEMENTS	50-134
50.75	FDRCLONE EXAMPLES	50-138
50.80	FDRDRP	50-141
50.81	FDRDRP JOB CONTROL REQUIREMENTS	50-147
50.82	FDRDRP RESTORE STATEMENT	50-148
50.83	FDRDRP SELECT STATEMENT	50-152
50.84	FDRDRP RESTORE EXAMPLES	50-154

TABLE OF CONTENTS

Page No.

51.0 FDRABR ARCHIVING AND SUPERSCRATCH

51.01	FDRABR ARCHIVING AND SUPERSCRATCH	51-1
51.02	DATA SET SELECTION	51-11
51.03	PROCESSING OPTIONS AND REQUIREMENTS	51-16
51.04	ARCHIVE AND SUPERSCRATCH JOB CONTROL REQUIREMENTS	51-22
51.05	ARCHIVE DUMP OR SCRATCH STATEMENT	51-28
51.06	SELECT, EXCLUDE AND MOUNT STATEMENTS	51-43
51.07	ARCHIVE RESTORE JOB CONTROL REQUIREMENTS	51-50
51.08	ARCHIVE RESTORE STATEMENT	51-52
51.09	SELECT STATEMENT FOR ARCHIVE RESTORE	51-60
51.20	ARCHIVE BACKUP EXAMPLES	51-68
51.21	AUTO RECALL AND POOLDISK EXAMPLES	51-70
51.22	ARCHIVE RESTORE EXAMPLES	51-72
51.23	SUPERSCRATCH EXAMPLES	51-74
51.24	SMS ARCHIVE EXAMPLES	51-75
51.30	ARCHIVE BACKUP ISPF SUPPORT	51-76
51.31	ABR AUTO-RECALL - INTRODUCTION	51-79
51.32	ABR AUTO-RECALL - IMPLEMENTATION	51-81
51.33	ABR AUTO-RECALL - OPERATION	51-83
51.34	ABR AUTO-RECALL - CUSTOMIZATION OPTIONS	51-88
51.35	ABR AUTO-RECALL - SECURITY	51-93
51.36	ABR AUTO-RECALL - HSM COMPATIBILITY	51-94
51.40	REMOTE QUEUE UTILITY (FDRABRUT)	51-95
51.41	FDRABRUT JOB CONTROL REQUIREMENTS	51-97
51.42	FDRABRUT ARCHIVE/BACKUP STATEMENT	51-98
51.43	FDRABRUT RESTORE STATEMENT	51-100
51.44	FDRABRUT RESET/REMOVE STATEMENT	51-101
51.45	FDRABRUT LIST STATEMENT	51-102
51.46	FDRABRUT EXAMPLES	51-103
51.50	ARCHIVE MAINTENANCE UTILITY (FDRARCH)	51-104
51.51	FDRARCH JCL REQUIREMENTS	51-106
51.52	FDRARCH FORMAT STATEMENT	51-108
51.53	FDRARCH DUMP STATEMENT	51-110
51.54	FDRARCH RESTORE STATEMENT	51-111
51.55	FDRARCH MERGE STATEMENT	51-112
51.56	FDRARCH REORG STATEMENT	51-114
51.57	FDRARCH DELETE STATEMENT	51-122
51.58	FDRARCH RESET STATEMENT	51-126
51.59	FDRARCH EXPIRE STATEMENT	51-127
51.60	FDRARCH MODIFY STATEMENT	51-128
51.61	FDRARCH RECATALOG STATEMENT	51-134
51.62	FDRARCH UPDATE STATEMENT	51-139
51.63	FDRARCH EXAMPLES	51-141

TABLE OF CONTENTS

Page No.

52.0 FDRAPPL APPLICATION BACKUP

52.01	FDRAPPL APPLICATION BACKUP	52-1
52.02	DATA SET SELECTION	52-7
52.03	PROCESSING OPTIONS AND REQUIREMENTS	52-9
52.04	FDRAPPL BACKUP JOB CONTROL REQUIREMENTS	52-13
52.05	FDRAPPL DUMP STATEMENT	52-16
52.06	SELECT STATEMENT FOR FDRAPPL	52-22
52.07	FDRAPPL RESTORE JOB CONTROL REQUIREMENTS	52-25
52.08	FDRAPPL RESTORE STATEMENT	52-27
52.09	APPLICATION RESTORE SELECT STATEMENT	52-33
52.10	MANAGING THE APPLICATION CONTROL FILE.....	52-37
52.11	USING A PERMANENT APPLICATION CONTROL FILE	52-38
52.12	USING A GDG APPLICATION CONTROL FILE	52-45
52.20	REPORTING ON APPLICATION BACKUPS	52-51
52.21	ISPF SUPPORT FOR APPLICATION BACKUPS	52-53

53.0 STANDARD REPORTING FACILITY (FDRABRP)

53.01	STANDARD REPORTING FACILITY (FDRABRP)	53-1
53.02	FDRABRP JCL REQUIREMENTS	53-3
53.03	FDRABRP ARCHIVE REPORT	53-4
53.04	FDRABRP CATALOG REPORT	53-8
53.05	FDRABRP SCRATCH REPORT	53-10
53.06	FDRABRP VOLUME STATUS REPORT	53-13
53.07	FDRABRP VTOC REPORT	53-16
53.08	FDRABRP DATA SET BACKUP REPORT	53-19
53.09	FDRABRP TAPE VTOC REPORT	53-22
53.10	FDRABRP JCL EXAMPLES	53-25

TABLE OF CONTENTS

Page No.

54.0	GENERALIZED REPORT WRITER (FDREPORT)	
54.01	GENERALIZED REPORT WRITER (FDREPORT)	54-1
54.02	INNOVATION HEALTH CHECK	54-4
54.03	FDREPORT JCL REQUIREMENTS	54-6
54.10	FDREPORT DEFAULT STATEMENT	54-9
54.11	FDREPORT TITLE STATEMENT	54-12
54.12	FDREPORT HEADING STATEMENT	54-13
54.13	FDREPORT (X)SELECT/(X)EXCLUDE STATEMENT	54-14
54.14	FDREPORT REPORT STATEMENT	54-20
54.15	FDREPORT SORT STATEMENT	54-21
54.16	FDREPORT BREAK STATEMENT	54-23
54.17	FDREPORT SUMMARY STATEMENT	54-24
54.18	FDREPORT PUNCH STATEMENT	54-25
54.19	FDREPORT PRINT STATEMENT	54-28
54.20	FDREPORT CANCEL STATEMENT	54-39
54.21	FDREPORT EXECUTE STATEMENT	54-40
54.30	FDREPORT EXECUTION OPTIONS	54-41
54.31	FDREPORT FIELD NAMES	54-45
54.32	FDREPORT HSM REPORTING	54-55
54.40	FDREPORT VTOC EXAMPLES	54-57
54.41	FDREPORT ARCHIVE EXAMPLES	54-66
54.42	FDREPORT CATALOG EXAMPLE	54-68
54.43	FDREPORT SMS EXAMPLES	54-69
54.44	FDREPORT VOLUME EXAMPLES	54-71
54.45	FDREPORT PUNCH EXAMPLES	54-73
54.46	FDREPORT TAPE EXAMPLES	54-77
54.47	FDREPORT YEAR 2000 EXAMPLES	54-79
54.48	FDREPORT HSM EXAMPLES	54-80
54.50	FDRSRS -- SEARCH, REPORT AND SERVICES DIALOG	54-81
54.51	SRS DIALOG OPTIONS AND DEFAULTS PANELS	54-85
54.52	SRS DATA SET SELECTION CRITERIA PANEL	54-90
54.53	SRS DATA SET LIST PANEL	54-95
54.54	SRS DATA SET LIST FDR/ABR FUNCTIONS	54-97
54.55	SRS VOLUME SELECTION CRITERIA PANEL	54-98
54.56	SRS VOLUME LIST PANEL	54-99
54.57	SRS ISPF FASTPATHS AND COMMANDS	54-100
54.60	SRS EXAMPLES	54-101

TABLE OF CONTENTS

Page No.

60.0	FDR AND ABR BACKUP MAINTENANCE	
60.01	FDR AND ABR BACKUP MAINTENANCE	60-1
60.10	FDR TAPE COPY (FDRTCOPY)	60-3
60.11	FDRTCOPY JCL REQUIREMENTS	60-6
60.12	FDRTCOPY CONTROL STATEMENTS	60-8
60.13	FDRTCOPY EXAMPLES	60-15
60.20	FDRTSEL INTRODUCTION	60-18
60.21	FDRTSEL JCL REQUIREMENTS	60-19
60.22	FDRTSEL CONTROL STATEMENTS	60-22
60.23	FDRTSEL COPY/MOVE STATEMENT	60-23
60.24	FDRTSEL SELECT ARCHIVE/ARCEDIT STATEMENT	60-24
60.25	FDRTSEL SELECT CATLG STATEMENT	60-29
60.26	FDRTSEL CONTROL STATEMENT	60-31
60.27	FDRTSEL FEATURES AND CONSIDERATIONS	60-33
60.28	FDRTSEL ARCHIVE EXAMPLES	60-36
60.29	FDRTSEL BACKUP EXAMPLES	60-39
60.30	FDRCOPY/FDRTSEL ISPF INTERFACE	60-42
70.0	SYSTEM MANAGED STORAGE (SMS)	
70.01	SYSTEM MANAGED STORAGE (SMS)	70-1
70.10	SMS VOLUME OPERATIONS	70-3
70.11	SMS DATA SET OPERATIONS	70-4
70.12	SMS ABR SUPPORT	70-7
70.13	SMS ACS ROUTINES	70-16
70.20	SMS QUESTIONS AND ANSWERS	70-17
70.30	FDRCONVT OVERVIEW	70-24
70.31	PROCESSING OPTIONS AND FEATURES	70-30
70.32	FDRCONVT JCL REQUIREMENTS	70-31
70.33	FDRCONVT CONVERTVOL STATEMENT	70-32
70.34	FDRCONVT SELECT STATEMENT	70-37
70.35	FDRCONVT EXAMPLES	70-38
80.0	SPECIAL CONSIDERATION OVERVIEW	
80.01	SPECIAL CONSIDERATION OVERVIEW	80-1
80.10	DATA SET RESTORE RULES	80-2
80.11	FDR PROCESSING BY TYPE OF DATA SET	80-7
80.12	DATA MOVEMENT BETWEEN DIFFERENT DASD DEVICES	80-15
80.13	ICF VSAM SPECIAL CONSIDERATIONS	80-18
80.14	GENERIC DATA SET NAME SELECTION	80-29
80.15	SECURITY	80-34
80.20	DISASTER RECOVERY	80-42
80.30	FDR/DSF/CPK/ABR I/O ERROR RECOVERY	80-52
80.31	ABR ABNORMAL TERMINATION PROCEDURES	80-55
80.32	ABR AND TAPE MANAGEMENT SYSTEMS	80-56
80.33	STORAGETEK PERFORMANCE PRODUCTS	80-62
80.90	FDRCHSM CONVERSION AID	80-67

TABLE OF CONTENTS

Page No.

90.0	FDR INSTALLATION INTRODUCTION	
90.01	FDR INSTALLATION INTRODUCTION	90-1
90.02	RE-INSTALLATION OF THE FDR DISTRIBUTION TAPE	90-5
90.03	LOADING THE FDR LIBRARIES FROM THE DISTRIBUTION TAPE	90-7
90.10	INVOKING THE ABR INSTALL ISPF DIALOG (PANEL A.I)	90-12
90.11	SELECT THE FDR GLOBAL OPTIONS (PANEL A.I.4)	90-14
90.12	SECURITY OPTIONS (PANEL A.I.4.1)	90-18
90.13	GENERAL OPTIONS (PANEL A.I.4.2)	90-19
90.14	COMPAKTOR OPTIONS (PANEL A.I.4.3)	90-21
90.15	ABR GENERAL OPTIONS (PANEL A.I.4.4)	90-23
90.16	ABR DATA SET NAMES (PANEL A.I.4.5)	90-26
90.17	ARCHIVE UTILITY DEFAULTS (PANEL A.I.4.6)	90-28
90.18	ABR DISK PROCESSING OPTIONS (PANEL A.I.4.7)	90-30
90.19	MORE ABR GENERAL OPTIONS (PANEL A.I.4.8)	90-32
90.20	REPORT OPTIONS (PANEL A.I.4.9)	90-33
90.21	MORE REPORT OPTIONS (PANEL A.I.4.10)	90-35
90.22	OPERATING SYSTEM EXIT OPTIONS (PANEL A.I.4.11)	90-36
90.23	SET FDRREORG OPTIONS (PANEL A.I.4.12)	90-44
90.24	ENABLE FDRCLONE OPTIONS (PANEL A.I.4.13)	90-52
90.30	DEFINE THE COMPAKTOR UNMOVABLE TABLE (PANEL A.I.5)	90-53
90.31	DEFINE THE FDRREORG NOREORG LIST (PANEL A.I.5A)	90-55
90.32	DEFINE THE ABR PROTECT LISTS AND RESTORE ALLOCATION LIST (PANEL A.I.6)	90-57
90.40	AUTHORIZE THE FDR PROGRAMS	90-62
90.41	CREATE THE ABR CATALOG (PANEL A.I.7)	90-64
90.42	SET ABR DISK VOLUME PROCESSING OPTIONS (PANEL A.I.8)	90-66
90.43	CREATE THE ARCHIVE CONTROL FILE (PANEL A.I.9)	90-71
90.44	SET FDR DIALOG GLOBAL OPTIONS (PANEL A.I.10)	90-73
90.45	CREATE THE REMOTE QUEUE FILES	90-76
90.46	INSTALLING THE FDR DIALOGS	90-77
90.47	LOAD SAR ONTO DISK	90-80
90.50	TESTING THE FDR SYSTEM	90-81
90.60	DYNAMIC INSTALLATION OF THE FDR EXITS	90-85
90.61	DYNAMIC EXIT INSTALLATION PROCEDURE	90-87
90.62	TESTING THE FDR MVS EXITS	90-91
90.70	APPLYING FDR FIXES (FDRSPZAP)	90-95
91.0	FDR ALTERNATE INSTALLATION PROCEDURE	
91.01	FDR ALTERNATE INSTALLATION PROCEDURE	91-1
91.02	LOADING THE FDR LIBRARIES	91-2
91.03	EXECUTING FDRZAPOP	91-3

THIS PAGE WAS INTENTIONALLY LEFT BLANK

100.0 INTRODUCTION TO MESSAGES

100.01	INTRODUCTION TO MESSAGES	100-1
100.02	FDR CONSOLE MESSAGES AND TSO TERMINAL MESSAGES	100-3
100.03	MESSAGES FROM FDR, DSF AND ABR (FDRNNN)	100-14
100.04	MESSAGES FROM SAR (FDR9NN)	100-89
100.05	MESSAGES FROM ALL PROGRAMS (FDR99N)	100-95
100.06	COMPAKTOR MESSAGES (CPKNNNX)	100-96
100.07	FDRREORG MESSAGES (FDRRNN/FDRSNN)	100-111
100.08	ABEND CODES	100-126
100.09	WAIT STATE CODES FROM SAR	100-129

01.01 ORGANIZATION OF THIS MANUAL

**FORMAT OF
THE MANUAL**

The user manual is divided into sections, by general topic, identified by an integer such as 10 or 21. Sections are divided into subsections by subtopic, with identifiers such as 10.03 or 52.12. Within subsections, subtitles may appear in the left margins to help you quickly find topics of interest. Each page displays its subsection title in the top center of the page, and the subsection number at the top right corner. Within each section, pages are numbered with the section number and a sequential page number within the section, e.g., page 80-12. The sections are distributed in two binders, for ease of use.

VOLUME 1

The first volume of the two-volume printed set contains:

	COVER AND TABLE OF CONTENTS
01	MANUAL ORGANIZATION AND SUMMARY OF MODIFICATIONS
02	INTRODUCTION AND OVERVIEW OF FDR COMPONENTS
10	FAST DUMP RESTORE (FDR)
15	STAND-ALONE RESTORE (SAR)
20	DATA SET FUNCTIONS (DSF)
21	FDRCOPY
25	FDR INSTANTBACKUP
30	FDRREORG
40	COMPAKTOR (CPK)
50	ABR VOLUME BACKUPS AND FDRCLONE
51	ABR ARCHIVE AND SUPERSCRATCH

VOLUME 2

The second volume of the two-volume printed set contains:

52	ABR APPLICATION BACKUPS (FDRAPPL)
53	ABR REPORTS (FDRABRP)
54	FDREPORT AND FDRSRS
60	FDRTCOPY AND FDRTSEL COPY UTILITIES
70	SYSTEM MANAGED STORAGE (SMS)
80	SPECIAL CONSIDERATIONS
90	INSTALLATION AND OPTIONS
100	MESSAGES AND CODES

**IMPROVING
THE MANUAL**

We have tried to make this manual as complete, precise and error free as possible. However, in spite of our best efforts, errors and incomplete explanations may have crept in. Should you encounter any of these, we would appreciate your corrective criticism. It is primarily through your feedback that we can improve this manual. There is a Reader's Comment Form at the back of the manual; you may mail it, fax it, or send us e-mail.

EXAMPLES

All examples and Job Control Language statements shown in this manual are for ILLUSTRATIVE PURPOSES ONLY! You must modify them to meet the JCL requirements of your installation.

All sample jobstreams in this manual are also included in a JCL library on the distribution tape. If you load this JCL library to disk, you can reduce the time required to setup FDR jobstreams by using the JCL library members as a starting place.

01.02 CONTROL STATEMENT FORMAT

Control statements for FDR components consist of 80-character logical records ("card images") where:

- COLUMNS 1 to 71 – Contain the command, operands, and comments fields, except when continued to subsequent logical records.
- COLUMN 72 – Must be blank if the last character of the command or operands is in column 71
- COLUMNS 73 TO 80 – Not used by FDR; may contain an identification or sequence number.

**GENERAL
FORMAT**

For most FDR programs, the general format of control statements is:

COMMAND OPERANDS COMMENTS

where:

**COMMAND
FIELD**

The COMMAND field identifies the control statement. Each FDR component has a set of commands it will accept, as documented in the rest of this manual.. The command can start in column 1 of the input record, or it can optionally be preceded by any number of blank columns, as long as it ends before column 72. It cannot be continued to another input record. It must be followed by at least one blank column.

**OPERAND
FIELD**

The OPERAND field, if present, follows the command field and is separated from it by at least one blank. The operand field consists of one or more operands, **separated by commas** (a common mistake is separating the operands by blanks instead of commas, causing the extra operands to be taken as comments) . The operand field may not contain embedded blanks except within quoted strings. Operand fields may be continued onto subsequent logical input records but the first operand must appear on the same record as the command (see "Continuing an Operand Field" below).

Most FDR commands require operands. The operands accepted with each command are documented in the rest of this manual.

Most operands are keyword, meaning that multiple operands may appear in any order. In a few cases, the operands must appear in a specific order; this will be documented under the specific command.

**COMMENTS
FIELD**

The comments field, if present, follows the operand field and is separated by one or more blank columns. It may contain any characters; it is not validated. Comments fields may not be continued, i.e., they must end at or before column 71. Comments are not permitted on a control statement that allows operands but on which no operands have been specified.

You can also specify an entire record of comments by placing an asterisk (*) in column 1 of an input record. You can use the rest of the record for comments.

SUBOPERAND

Some operands consist of a list of suboperands. A suboperand list must be enclosed within parentheses, unless the list reduces to a single suboperand, in which case the parentheses may be omitted.

**QUOTED
STRINGS**

When the variable data you specify for a parameter contains certain special characters, defined below, you must enclose the data with apostrophes. This is called a 'quoted string'. Within a quoted string, all characters, including blanks, may appear; if an apostrophe is to be part of the string, it must be coded as two apostrophes. Quoted strings are accepted only where the documentation for an operand indicates so.

SPECIAL CHARACTERS: Parentheses, commas, equal signs, apostrophes, and blanks.

CONTROL STATEMENT FORMAT

CONTINUED

CONTINUING AN OPERAND FIELD

When the total length of the operands on a statement exceeds the available columns in a logical record, they must be continued onto one or more following logical records.

To continue a statement, interrupt the operand field after any complete operand or suboperand, including the comma which follows it. That comma must occur at or before column 71 and the next column must be blank. On the next input record, start the next operand anywhere in the record (columns 1 to 71).

You may continue the control statement any number of times. Some users put only one operand per record to improve readability.

SAMPLE CONTROL STATEMENTS

Here are some examples of properly formatted control statements, including continuation and comments:

```
COLUMN 1
|
DUMP TYPE=FDR

      RESTORE      TYPE=DSF ,DATA=USED      RESTORE DATA SETS

              DUMP      TYPE=APPL ,      APPLICATION BACKUP
              ARCBACKUP=NO ,      DO NOT BACKUP CONTROL FILE
              PRINT=DSF      LIST SELECTED DATA SETS

*      THE FOLLOWING IS A RESTORE STATEMENT - THIS IS A COMMENT
      RESTORE      TYPE=ARC ,
RECAT=YES

*      THE FOLLOWING ILLUSTRATES CONTINUING A SUBOPERAND
      SELECT  DSN=SYS1 . PARMLIB ,NEWNAME=TEMP . PARMLIB ,NVOL=( TEMP01 ,
              TEMP02 ,TEMP03 )
```

NOTATION

Each subsection of this manual which describes a control statement starts with a table showing the command and operands for that command, followed by detailed descriptions of the operands. In that table, the following notation is used in this manual to define the usage of each operand:

- Uppercase letters and words must be coded exactly as shown; they are also printed in bold.
- Lowercase non-bold letters and words represent variables for which you must substitute specific information as defined in the detailed descriptions
- When an operand may have several values, they are separated by a vertical bar. You must choose one of the values. For example, RECALL=YES|NO means that you can code RECALL=YES or RECALL=NO
- When a particular value is the default for an operand, it is underlined. For example, DATA=ALL|USED indicates that DATA=USED is the default. However, for some operands the default that is distributed with FDR can be changed by the installation by setting options in the FDR Global Option Table. In that case, to avoid misleading users who may not be aware that defaults were changed, the distributed default is not underlined; the descriptive text will identify this situation

It is strongly recommended that you specify all relevant options on each control statement, rather than depending on the setting in the FDR option table. This will insure proper execution of your job even if the option table is changed or reset, and will help Innovation diagnose problems in your FDR jobs.

- Some operands cannot appear together on a given statement. These and other restrictions are given in the detailed operand descriptions.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

01.03 YEAR 2000 COMPLIANCE

Innovation Data Processing considers its products to be Year 2000 (Y2K) compliant if the product is capable of correctly processing dates within the 20th and 21st centuries, including proper processing of leap years, provided that the products are used on platforms (hardware, operating systems, and associated products from other vendors) that are also Year 2000 compliant.

Dates that are used by Innovation Data Processing's products fall into three categories:

- Those which are part of system data areas such as the last referenced dates, expiration dates, and creation dates in the VTOC, VVDS, catalog, and tape labels, as well as the current date (run date). These dates are handled by Innovation's products in the same fashion as the operating system. Most of these dates are valid through the year 2155. Special dates, such as 99.365 or 99.366 for permanent retention, or 99.000 used for catalog control by some tape management systems, will be properly handled.
- Dates or durations (number of days) specified on control statements. Most date parameters have been updated to accept a 2-digit or 4-digit year; for 2-digit years any year number less than 70 is assumed to be in the 21st century (20xx). Parameters which specify a duration (such as ADAYS=, CRDAYS=, XDAYS=) calculate a future date or a prior date and will properly handle Y2K and leap years.
- Cosmetic dates, such as those which appear in printed output. Where possible, 4-digit years are used in reports (4-digit years are now the default in FDREPORT but can be overridden). Most dates in heading lines and similar displays use 4-digit years. In some cases, these dates may still appear with two digit years.

**FDREPORT
CUSTOMERS**

The default in FDREPORT is now DATEFORMAT=YYYYDDD. By default, dates are displayed with a 4-digit year and the line length of some existing reports with dates may expand (compared to V5.2) and become too long to display. If you prefer, you can revert to 2-digit years by putting DATEFORMAT=YYDDD on a DEFAULT or PRINT statement. [See Section 54](#) for instructions on permanently changing the DATEFORMAT= default for all FDREPORT jobs.

For expiration dates of 99000, 99365, 99366 or 99999 (the latter for VSAM only), FDREPORT will set EXPDAYS or BKEXDAYS to 65535.

CONTINUED. . .

YEAR 2000 COMPLIANCE

CONTINUED

EXPIRATION DATES

If you code EXPDT= rather than RETPD= on TAPEX or TAPEXX DD statements used with ABR, then:

- In the backup subsystem (full-volume and incremental backups), ABR will interpret any EXPDT value as a keyword expiration.
- In the ARCHIVE subsystem (ARCHIVE and application backup), ABR will interpret the date as a keyword expiration only if you use the values:
 - EXPDT=99000 - indicates catalog control to most tape management systems
 - EXPDT=99365 or 99366 (tape or disk output) - indicates permanent retention.

The term "keyword expirations" refers to expiration dates that have special meanings to tape management systems and to IBM data management. ABR is coded to recognize and properly handle keyword expirations as indicated above. For ARCHIVE tapes, 99000 and 99365 are the only keyword expirations recommended for use with ABR; ABR may not properly handle other keyword expirations (such as 98ddd for "retain until unused for ddd days"). If you are currently using any keyword expiration for ARCHIVE tapes other than the recommended values, **please contact Innovation for assistance.**

Many tape management systems include options to treat keyword expiration dates as real dates. For example:

```
//REALDATE DD DUMMY      for CA-DYNAM/TLMS
//TMNOKEY  DD DUMMY      for CA-1
TRUXPD=YES                global option for CA-1
```

ABR will not recognize these options and will use its normal rules for treating EXPDT values as keywords. If you are using these options in ABR jobs, **please contact Innovation for assistance.**

You may use EXPDT= to specify expiration dates that are real dates and not keyword expirations (e.g., EXPDT=99304 or EXPDT=2001/123). ABR full-volume and incremental backups will treat such dates as keyword expirations, possibly giving incorrect results. ARCHIVE and application backups will correctly treat such dates as real dates.

However, it is rare that explicit expiration dates are used with ABR. It is more common to use retention periods (the RETPD= parameter on the TAPEX or TAPEXX DD statement or RETPD= on the ABR DUMP statement or a default retention period provided by ABR). ABR will properly handle calculation of the expiration date from a retention period, including all leap year and Y2K considerations.

CONTINUED . . .

01.04 SUMMARY OF MODIFICATIONS FOR VERSION 5.3

SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 50

ABR INSTANTBACKUP SUPPORT FOR HDS SHADOWIMAGE	Customers who are licensed both for FDR/ABR and for FDR InstantBackup for Hitachi Data Systems ShadowImage (on the HDS 7700E disk subsystem) can now use ABR to do non-disruptive point-in-time backups of volumes with ShadowImage.
INSTANTBACKUP SUPPORT FOR FLASHCOPY	FDR InstantBackup now supports the FlashCopy feature on the IBM 2105 ESS "Shark" disk subsystem. Customers who have licensed FDR InstantBackup for FlashCopy can now use FDR and FDRDSF to do non-disruptive point-in-time backups of Flashcopy-capable disks. Customers who are licensed for both the FDR/ABR and for FDR InstantBackup for FlashCopy can now use ABR to do non-disruptive point-in-time backups of volumes with FlashCopy.
ABR INSTANTBACKUP UP EMC SYMMETRIX SPLITS	The ABR SPLIT command for EMC Symmetrix Timefinder will automatically use "differential split" on microcode level 5x65 whenever the BCV being split has its own M2 mirror; this will improve the performance of future SPLITS since only updated tracks will need to be copied to the mirror. On microcode level 5x66, ABR SPLIT will automatically use "instant split" which will improve the performance of all SPLITS, especially if the previous ESTABLISH has not yet completely synchronized the volumes; instant splits are also differential splits if a BCV has a mirror.
SUPPORT FOR MULTI- VOLUME DATA SETS	FDRDSF, FDRCOPY and FDRABR now support restore, copy, and move of multi-volume data sets residing on up to 59 disk volumes.
ALTERNATE INDEX RESTORE	<p>ABR can now restore VSAM clusters with alternate indexes (AIXs) on separate volumes, regardless of the order in which the volumes are restored. If ABR attempts to restore an AIX before its associated base cluster has been completely restored, ABR will bypass it and will automatically re-attempt the restore of the failed AIXs after all other requested restores are complete (this may involve remounting a backup tape). All types of ABR restores are supported, and it works even when the base cluster is multi-volume.</p> <p>ABR ARCHIVE can now archive and auto-recall VSAM clusters with alternate indexes (AIXs), even when the DD statement or dynamic allocation specifies a PATH name or AIX cluster name.</p>
FDRREORG SUPPORT FOR IAM AIX	FDRREORG can now reorganize an IAM data set with an alternate index. Unlike native VSAM AIXs, FDRREORG does not need to rebuild the IAM AIX when the base data set is reorganized.
NOTE FOR CA-DYNAM/ TLMSTAPE MANAGEMENT USERS	V5.5 of the CA-DYNAM/TLMS tape management system supports 3 new retention types (A, B, and C) which are useful for the retention of ABR-created tapes. These retention types are equivalent to types 1, 2, and 3 (for catalog control, date control, and catalog/date control) except that they will retain the tape until all files on the tape have expired. When these retention types are used, the TLMSXTRS exit modification previously used by ABR customers is no longer required. This applies to all levels of ABR, not just 5.3/50.

01.04 CONTINUED

SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 40

**SUPPORT FOR
THE HSDM
OPTION FROM
STORAGETEK**

V5.3/40 and above supports the HSDM (High Speed Data Mover) option of the SVA (Shared Virtual Array) disk subsystem from StorageTek; HSDM is also available on some models of the IBM RVA (Ramac Virtual Array). Contact StorageTek for information on the requirements and pricing on HSDM or visit their web site at:

www.storagetek.com

In an SVA/RVA, all data is compressed, even in cache. It is uncompressed only when it must be sent to the host. When backing up data in an SVA/RVA, it must be decompressed by the control unit, and recompressed when it is restored.

HSDM allows FDR to directly read the compressed track images, avoiding the decompression overhead. The amount of data that must be read from disk and written to the backup is reduced by 3:1 or more, which reduces the elapsed time of the FDR backup up to 60%. When restoring the same data to an SVA/RVA with the HSDM option, HSDM allows FDR to write the compressed track images directly, avoiding compression overhead and providing similar reductions in restore elapsed time.

When data dumped with HSDM must be restored to a non-HSDM disk, FDR calls a software decompression subroutine provided by StorageTek, and restores the uncompressed data to the output disk normally.

HSDM is supported by all FDR backup/restore functions in programs FDR, FDRDSF, and FDRABR as well as restore functions in FDRCLONE and FDRDRP.

To invoke HSDM during a backup, add the operand DCT=YES (DUMPCOMPRESSEDTRACK=YES) to the DUMP statement under any of the programs listed above; HSDM will be invoked for all volumes with HSDM enabled and will be ignored for other disks which will be dumped normally. No special parameters are required to restore a HSDM-created backup; FDR (including SAR) will automatically recognize the HSDM format and restore it properly.

You can use any of your existing FDR, FDRDSF, or FDRABR jobstreams, or any of the examples in the FDR manual, and simply add DCT=YES, for example,

```
DUMP TYPE=FDR,DCT=YES
DUMP TYPE=ARC,ADAYS=15,DSNENQ=USE,DCT=YES
```

In a test run at Innovation, dumping one 3390-3 volume from a SVA to a StorageTek 9840 tape, HSDM reduced elapsed time by 54%.

In order to use HSDM, you must be licensed for FDR InstantBackup and you must have the HSDM feature installed on your SVA or RVA.

More details on HSDM support can be found in [Section 80.33](#) and in the sections for each FDR function.

WARNING: if you use DCT=YES and have one of these IBM PTFs applied to your system

UW90591 UW90593 UW61573 UW90595 UW61575

then you must apply the PTF which fixes IBM APAR OW42799 to avoid a problem which may cause backups to fail. Read the description of APAR OW42799 for details.

CONTINUED . . .

01.04 CONTINUED

**ABR SUPPORT
FOR
STORAGETEK
EXHPDM**

ABR now supports the ExHPDM (Extended High Performance Data Mover) software product from StorageTek; in 5.3/30 it was supported only for FDR and FDRDSF, but that support is extended to ABR Volume Backups and Application Backups (FDRAPPL) in V5.3/40. It is not supported for ABR Archive Backups. ExHPDM interleaves several concurrent output data streams (such as concurrent FDR backups) onto one tape to improve performance and reduce the number of tape volume required.

ExHPDM is invoked by the SUBSYS= JCL parameter on TAPEX and TAPEXX DD statements in FDR jobsteps, e.g.,

```
// TAPE1 DD DSN=BACKUP.VPROD01(+1),UNIT=TAPE,
// SUBSYS=(SOV,'CLASS(FDRBKUP)'),DISP=(,CATLG)
```

More details on ExHPDM support can be found in [Section 80.33](#) and in the sections for each FDR function.

You must be licensed for FDR InstantBackup to use ExHPDM with ABR. You must be running at least ExHPDM V1.1 plus PTF L1P013W (for FDR support) and PTF L1P015A (for ABR support).

**IMPROVED
BACKUP
PERFORMANCE
AND
CONCURRENCY**

In previous releases, a backup to a sequential disk data set or to disk and tape concurrently would take much longer than a backup to tape alone, due to the disk I/O technique used on the input when outputting to disk. Backup to disk is most commonly used with ABR Archive (DUMP TYPE=ARC) to create COPY 1 on disk for rapid recall.

Also, each concurrent backup executing in an FDR step required about 1M of memory below the 16MB line, limiting the number of such concurrent backups to the size of the available below the line private area in each installation (usually 6-10MB).

In V5.3/40 and above, these limitations are removed. On any DUMP statement in FDR, FDRDSF or FDRABR, add the operand RTC=YES (READTRACKCCW=YES). With RTC=YES, most of the FDR buffers are moved above the 16MB line, allowing more concurrent backups to run, up to the limit of 9 in one address space. FDR will read up to 1 cylinder (15 tracks) of data in one I/O, compared to 1/2 cylinder without RTC=YES. Backups to disk will run much faster, and backups to tape will see some improvement depending on the data rate of your tapes and disks.

The backup produced by RTC=YES will have the same format as backups from previous releases and can be restored by any FDR release.

Note: in V5.3/40, COMPRESS= will be ignored if RTC=YES is specified, so RTC backups to disk will not be compressed. They may occupy more disk space than compressed backups without RTC=YES. Innovation intends to lift this restriction in a future release.

WARNING: if you use RTC=YES and have one of these IBM PTFs applied to your system

UW90591 UW90593 UW61573 UW90595 UW61575

then you must apply the PTF which fixes IBM APAR OW42799 to avoid a problem which may cause backups to fail. Read the description of APAR OW42799 for details.

**MESSAGE
ADR299E
RESOLVED**

When previous releases of FDR were used to do a full-volume restore or copy, and the IBM product DFSMSdss was later run on the volume, DFSMSdss might generate message ADR299E indicating that the volume size in the VTOC index was incorrect. FDR will now correct the VTOC index volume size during the restore/copy; the correction will occur only if both the backup and restore are done with V5.3/40 or above.

**ENHANCED
SUPPORT FOR
HFS DATA SETS
FOR UNIX
SYSTEM
SERVICES**

In order to insure that FDR can get consistent, usable backups of HFS (Hierarchical File System for OS/390 Unix System Services) data sets, all FDR, FDRDSF, and FDRABR backups now support the operand HFS=QUIESCE. If specified, and FDR is unable to get a SYSDSN ENQ on the HFS data set (which usually means that the file system is mounted to OS/390 Unix), FDR will issue a OS/390 Unix "quiesce" function which will suspend all access to the file system until the backup is complete. This applies to most FDR and FDRABR backups as well as FDRCOPY COPY and FDR InstantBackup SNAP and SPLIT; it does not apply to COMPAKTOR, Archive Backups or FDRCOPY MOVE.

FDR can backup and restore entire HFS data sets. If you need the ability to backup and/or recover individual HFS files within the data sets, this is provided by UPSTREAM for Unix Systems Services, a separately-priced product from Innovation.

01.04 CONTINUED

SUPPORT FOR NON-SMS PDSE AND HFS DATA SETS	<p>IBM has released PTFs UW65973 (DFSMS 1.4) and UW65974 (DFSMS1.5) for (OS/390 2.4 and above to support PDSE (PDS Extended) and HFS (Hierarchical File System for OS/390 Unix) data sets on non-SMS volumes. Previously they were supported only on SMS-managed volumes.</p> <p>FDR 5.3/40 or above must be installed if you plan to use non-SMS PDSE or HFS data sets.</p> <p>Toleration zaps are available for previous releases but they do not support allocation of these data sets.</p>
ENHANCED NVOL SUPPORT	<p>The NVOL= operand, used in DSF and ABR data set restores and FDRCOPY copy/move, has been enhanced. For multi-volume data sets it now allows you to more closely control the order in which the volumes specified by NVOL= will be used when allocating the data set. For details see the NVOL= description in various sections of this manual.</p>
MULTI-VOL VSAM ALLOCATIONS	<p>The technique used to allocate multi-volume VSAM clusters has been changed to improve the chance of a successful allocation when the target volumes are fragmented. Previously it was possible that FDR could allocate the first extent of a multi-volume cluster on a volume but would fail when trying to extend the cluster to additional extents (usually due to lack of free space). Now it will retry the allocation on another volume and report the error only if the allocation fails on all target volumes.</p>
FDR SUMMARY PRINTOUTS	<p>Many FDR backups and restores now support a FDRSUMM DD statement. This is usually a SYSOUT data set, and will contain a one line summary for each disk volume processed, showing result codes and statistics, such as elapsed time, track counts, and byte counts. When FDR gives an abend or return code indicating that a problem occurred, this is a convenient way to find the failing volume quickly.</p> <p>It is supported on all backups but only if RTC=YES or DTC=YES is specified on the DUMP statement.</p> <p>It is supported only for FDR, ABR, and FDRDRP full-volume restores.</p>
FDREPORT CATALOG PROCESSOR	<p>FDREPORT contains a new catalog processor. It makes all catalog fields available for reporting. Previously, some fields (such as creation/expiration date and SMS classes) were taken only from the VTOC/VVDS, not the catalog. Now they will be taken from the catalog as well. It is invoked for DATATYPE=CATALOG, CATVTOC, or CATARCH.</p>
FDREPORT EXTENDED FORMAT VSAM SUPPORT	<p>FDREPORT now supports Extended Format (EF) VSAM clusters. The size of clusters over 4GB are properly reported. For compressed clusters, the size of the cluster before and after compression can be reported.</p>
IBM 2105 SHARK	<p>All versions of FDR support the IBM 2105 SHARK disk subsystem without modifications.</p>
IBM 3590E CARTRIDGE DRIVES	<p>All releases of FDR support the IBM Magstar 3590E drives, as well as the new extended length cartridges.</p>

01.04 CONTINUED

SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 30

**FDR INSTANT
BACKUP FOR
ABR**

For customers licensed for FDR InstantBackup **and** FDR/ABR, ABR has been enhanced to use hardware facilities of several disk subsystems to provide "instant" volume backups of disks in those subsystems; volume backups include full-volume and incremental backups. Supported hardware facilities are:

- the optional Snapshot feature of the IBM RVA (Ramac Virtual Array) or StorageTek Iceberg or SVA (Shared Virtual Array) disk storage subsystems
- the optional Timefinder feature of the EMC Symmetrix ESP disk storage subsystem

When the proper facilities are available, a special FDRABR step can be executed to instantly create frozen full-volume images of online disk volumes; these images remain offline to the host. The offline images are actually backups, ABR can restore from them if necessary. The creation of the frozen image takes only seconds, so update activity on the online volumes needs to be quiesced for only the time necessary to create the images.

Once the frozen images have been captured, you can execute another FDRABR step with special parameters, to locate and backup the frozen image to tape in normal ABR fashion.

The end result is that the actual ABR full-volume or incremental backup completes in just seconds per volume, while the job of moving that backup to tape takes place asynchronously, when you have the time and resources to do it.

**FDRCLONE –
NEW
PRODUCT**

Customers who use ABR to perform ABR volume backups (full-volume and incremental backups) can license FDRCLONE, which provides the ability to quickly "clone" data from those backups onto a test system (such as a Y2K test system or LPAR) or restore at a disaster/recovery site.

Rather than massive restores, FDRCLONE lets you select the volumes or data sets to be "cloned" but will restore them only when a batch job or TSO user references them, using a process similar to ABR auto-recall of archived data sets. So, only the data sets which are actually needed are restored; this may be only a fraction of the actual data on the backups.

You can direct the restores to new volumes, so the volumes on the test/disaster do not have to have the same volume serials as the original volumes. Alternately, those systems can have common volumes but have fewer volumes on the test system.

FDRDRP

FDRCLONE also includes FDRDRP, a utility for optimizing full-volume recovery from ABR volume backups (**full-volume and incremental backups**). This can be used for volumes where full-volume recovery is more appropriate, such as system volumes and high-priority production volumes which must be recovered quickly.

A normal ABR full-volume recovery restores one disk volume at a time. Since a given backup tape may contain the backups of many disk volumes, the same tapes may be unloaded and remounted over and over, taking considerable time and overwhelming tape librarians and automated tape libraries.

FDRDRP processes multiple full-volume recovery tasks in parallel. It manages usage of the backup tapes required for those restores, so that each backup tape is mounted a minimum number of times, **usually one mount per tape volume**. This will greatly reduce the elapsed time required to recover the volumes and eliminate most extra tape mounts. Elapsed time savings are often as much as 80%.

**IMPROVE-
MENTS IN
RESTORE OF
ALTERNATE
INDEXES**

When clusters with alternate indexes (AIXs) are backed up, FDR will record the description of the VSAM PATHs associated with each AIX. During restore, if an AIX must be defined, its paths will also be defined. This allows clusters with alternate indexes to be immediately usable without requiring manual intervention.

Both the backup and the restore must be done with V5.3 level 30 or above.

01.04 CONTINUED

**ALIASES FOR
USER
CATALOGS**

When ICF VSAM user catalogs are backed up, FDR will record the names of all aliases associated with each catalog in the master catalog. During restore, if a user catalog must be defined, its aliases will also be defined in the master catalog. This allows user catalogs to be immediately usable without requiring manual intervention. Both single-level and multi-level aliases are supported.

You can bypass the definition of the aliases during allocation of the catalog by adding the operand "ALIAS=NO" to the RESTORE statement.

Both the backup and the restore must be done with V5.3 level 30 or above.

**CONCURRENT
RESTORES**

FDR can now restore multiple full-volume backups concurrently in a single step. If you specify the operand MAXTASKS=nn on the FDR RESTORE statement, it will execute up to that many restores concurrently. The job step must include multiple TAPEX/DISKx/SYSPRINx DD sets. TAPEX DDs which use the same tape drive (because of UNIT=AFF= or volume affinity) will not be done concurrently; they will wait until another restore task using the same drive terminates.

**FULL-VOLUME
ABR RESTORES
WITH
DYNTAPE**

ABR full-volume restore with DYNTAPE has been enhanced to retain the most recently mounted input tape when restoring multiple volumes in the same step. If the tape volume already mounted (for the last disk volume restored) is also required for the next disk volume, it will be used without dismounting and remounting the tape. If a different tape volume is required, the previous tape will be deallocated and the new tape will be allocated (which may involve a different tape drive).

This is in support of high-capacity tape cartridges (such as IBM Magstar and StorageTek Redwood and 9840) which may contain backups of many disk volumes. This enhancement in V5.3 level 30 supersedes the suggestion made in earlier V5.3 manuals to use a TAPEX DD instead of DYNTAPE in this environment.

**ABR SORTS
DISK VOLUMES**

Whenever you execute ABR with more than one TAPEX DD statement, it backs up disk volumes in parallel. Now ABR will sort the disk volumes by the last digit of their MVS device address to attempt to spread the FDR workload across channels and control units to reduce contention. This only applies to volumes selected by the ONLINE or ONLVOL operands and MOUNT statements; volumes specified by DISKxxxx DD statements will continue to be processed first, in the order that the DDs appear. The VOLSORT= operand controls this function, but it defaults to YES, so if you have multiple TAPEX DDs, you may see volumes processed in a different order than in previous releases.

**StorageTek
9840**

FDR and ABR support the StorageTek 9840 high-capacity cartridge drive (20GB before compression). As with any high-capacity cartridge, ABR will automatically stack files on the tape and take advantage of the capacity of the tape.

**StorageTek
ExHPDM**

Program FDR and FDRDSF support ExHPDM, a software product from StorageTek which combines multiple FDR output files into one tape file, maximizing tape utilization and efficiency. V5.3 level 30 or higher is required for use with ExHPDM. [Section 80.33](#) has more details on ExHPDM usage with FDR. Contact StorageTek for details on ExHPDM or visit their web site at www.storagetek.com

**INTERNET
SALES AND
SUPPORT**

The Innovation site on the World Wide Web (WWW) contains the latest information on Innovation products. From that site you can easily determine latest release of the product, as well as the release you must be executing to support certain software (e.g., OS/390) and hardware (e.g., RAMAC) products before you upgrade.

Since most users now have easy access to the Internet, an increasing number of technical questions are being sent by e-mail. Our response time for responding to a technical question sent by e-mail is usually a few hours after receipt during normal business hours. You should always provide your name, your company name, and phone number in the message. Our Web site has links allowing you to send us e-mail, or you can use the direct e-mail address on the cover page of this manual. **URGENT PROBLEMS SHOULD ALWAYS BE CALLED IN FOR AN IMMEDIATE RESPONSE.**

The Innovation Web site can also be used to order additional marketing information and new releases of products for which you are licensed. When ordering new releases, you need to provide us with your complete mailing address and phone number, plus the location of the CPU if it is different from your mailing address.

EUROPEAN USERS: see the cover page of this manual for European office mailing and e-mail addresses and phones.

CONTINUED . . .

01.04 CONTINUED

**SPECIAL
CONSIDERA-
TIONS FOR
USERS OF
PRIOR
VERSIONS**

Please review [Section 01.03 Year 2000 Compliance](#). Depending on the FDR components included in your license, you may need to be at specific levels of FDR in the year 2000.

If you were given custom zaps to FDR components for an earlier release, do not attempt to apply them to V5.3. Contact Innovation for support.

Remember that backup tapes created by FDR components cannot be copied by normal copy programs such as IEBGENER. The copy program may appear to execute normally, but it will corrupt the backup on the output tape. You should always use the FDR utilities FDRTCOPY or FDRTSEL ([see Section 60](#)) to copy FDR backups.

**SOFTCOPY
DOCUMENT-
ATION**

User manuals for all Innovation products are available on a CD-ROM in two formats:

- IBM Bookmanager format for online viewing and searching. These can be viewed on your workstation or uploaded to MVS for viewing with Bookmanager READ/MVS (included in OS/390). If you don't have a licensed copy of Bookmanager READ on your workstation, the IBM Library Reader (a limited copy of Bookmanager READ) is also on the CD or you can use the Library Reader from the IBM documentation CD-ROMs.
- ADOBE PDF files for printing. Using ADOBE ACROBAT READER on your workstation (also included on the CD), you can view and print manuals that look *exactly* like Innovation's printed manuals.

One copy of the CD-ROM is available to each licensed installation at no additional charge.

**NEWS VIA
E-MAIL**

You can now receive news, announcements, new product information, latest versions and technical bulletins via E-mail. You can subscribe to this mailing list at our web site www.innovationdp.fdr.com/list.htm

**PRODUCT
NEWS**

Member NEWS in the ICL (Installation Control Library, loaded as part of FDR installation) contains the latest news and notes about the FDR family of products. It may contain details on changes made too late to be included in this manual. **All users should review the NEWS member after installation.**

SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 22**FDR INSTANT
BACKUP**

For customers licensed for FDR InstantBackup, FDRCOPY has been enhanced to automatically use special hardware facilities to copy tracks when it detects that they are available:

- when it determines that the input and output volume for a given data set are in the same IBM RVA or StorageTek Iceberg or SVA subsystem with the Snapshot feature, Snapshot commands will be used to provide an almost instantaneous copy or move.
- when it determines that the input and output volume for a given data set are in the same EMC Symmetrix subsystem, internal track copy commands will be used to reduce the copy time since the track data is not sent over the channel.
- Standard read/write commands will be used if the input and output volumes are not in the same RVA/Iceberg/SVA or Symmetrix.

This function is automatic, no special options are required to use it.

CONTINUED . . .

01.04 CONTINUED

SUMMARY OF MODIFICATIONS FOR V5.3 LEVEL 20

**CHANGES
MADE TO ALL
PRODUCTS**

The FDR interactive install is now available to customers licensed for any or all of the FDR components; previously it was available only for ABR customers. The interactive install is documented in [Section 90](#).

The installation process now loads a JCL library from the distribution tape. The JCL library contains all example JCL which appears in this manual. Member names are EXssssa where "ssss" is the section number (e.g., 5110 for Section 51.10) and "a" is a sequence character (e.g., A for the first example in that section).

**FDR INSTANT
BACKUP**

FDR InstantBackup is a new, extra-cost feature of FDR. It allows you to make use of features on
 EMC Symmetrix (Timefinder)
 IBM 9393 RVA (Ramac Virtual Array) (Snapshot)
 StorageTek Iceberg or SVA (Shared Virtual Array) (Snapshot)
 HDS 7700/7700E (Data-Plex HCPF and Shadow Image HMRCF)
 COMPAREX T2000/T2100 (CDRF/CMRF)

which use hardware features to make a duplicate of one online volume to another volume that is offline. On each vendor's subsystem, the offline copy can be frozen at a point-in-time even though updates continue on the online volume. Once the copy is frozen, FDR InstantBackup allows you to backup the offline copy with FDR or FDRDSF, or to copy data sets from the offline volume with FDRCOPY.

FDR InstantBackup, in conjunction with the various vendors' features, enables you to create a point-in-time backup of an online volume, while quiescing updates to the online volume for only the few seconds necessary to freeze the copy.

COMPAKTOR customers who are also licensed for FDR InstantBackup can use the hardware features of EMC Symmetrix and IBM RVA disk subsystems to move data on the disk without sending the data to and from the CPU. Such COMPAKTions will usually run faster.

ABR support for FDR InstantBackup on EMC Symmetrix, IBM RVA and StorageTek Iceberg/SVA is available in V5.3 level 30.

Note: customers who use FDR InstantBackup V5.2 with IBM RVA or StorageTek Iceberg or SVA DASD should review [Section 26](#). FDR InstantBackup has been enhanced to issue Snapshot commands directly, greatly simplifying the job of managing the Snapshot target volumes.

**CHANGES
MADE TO
COMPAKTOR**

The default values of SIZEKEEP have changed. In Version 5.3 and beyond the default is SIZEKEEP=(100,90,60). Please read the descriptions of SIZEKEEP in [Section 40](#) to understand the implications. Innovation strongly recommends that you run with the default values for most COMPAKTOR runs.

Users who execute COMPAKTOR under PGM=FDR or PGM=FDRABR (the COMPAKT operand on the DUMP statement) should change those jobs to do a separate Fast COMPAKTion step (PGM=FDRCPK with TYPE=FASTCPK) after the FDR or ABR step.

**CHANGES
MADE TO ABR**

In support of the newer high-capacity tape cartridge drives, such as the IBM Magstar and the StorageTek Redwood, ABR now supports up to 4095 backup files on a backup tape or multi-volume tape set; the previous limit was 255 which is still the default.

In addition, restore from Archive Backups will support the high-speed locate function for all types of cartridge drives. This allows ABR to position directly to the beginning of a backup file at high speed without reading the labels of all the files in-between. This will allow for faster recalls, especially when only one or two data sets are being restored. The positioning information is stored in the Archive Control File but the format of that file is unchanged and is compatible with V5.2 and earlier releases.

Restores from Volume Backups will also use high-speed locate, but only on IBM Magstar tapes. If the LASTAPE option is used for any kind of ABR backup, high-speed locate is used on Magstar tapes when positioning to the LASTAPE file in order to add files to an existing backup tape.

CONTINUED . . .

01.04 CONTINUED

**CHANGES
MADE TO
FDREPORT**

FDREPORT can now extract and report on data from the BCDS and MCDS control files of DFSMSHsm (or DFHSM). So, FDREPORT can be used as a reporting tool by HSM users. It can report on data sets migrated by HSM (MCDS) and data sets with backups under HSM (BCDS). All the power of the FDREPORT selection criteria and formatting can be used to customize the report.

The default in FDREPORT is now DATEFORMAT=YYYYDDD. Dates are displayed with a 4-digit year unless overridden by the user ([see Section 01.03](#)).

A new option, ENABLE=CHAREXPDATES is available. If this option is in effect, EXPDATE or BKEXDATE will be displayed as "CATCTL" for an expiration date of 1999.000 (99000) or "NEVER" for expiration dates of 1999.365, 1999.366, or 1999.999.

**NOTES FOR
ALL FDR AND
ABR USERS**

Performance: if you are doing backups with COMPRESS= on the DUMP statement and directing the backups to:

- IBM 3590 Magstar
- IBM 3490E on ESCON channels
- StorageTek Timberline
- StorageTek Redwood
- StorageTek 9840

Innovation strongly recommends that you remove COMPRESS= and use the hardware compaction features (e.g., IDRC) of those tape drives. Hardware compaction is requested by DCB=TRTCH=COMP on the TAPEX DD statement, but it may also be the default in your system.

Removing COMPRESS= also reduces the memory requirements of FDR and ABR by 50%, so you may want to add more TAPEX DD statements to do more concurrent backups.

MIGRAT=YES

The MIGRAT=YES option of Archive catalogs archived data sets so that the first or only volume serial is MIGRAT, which is recognized by IBM as indicating an archived data set. MIGRAT=YES is required for SMS-managed data sets.

Innovation recommends that all customers who are not yet using MIGRAT=YES for non-SMS archives should convert to using it when installing V5.3. So that you do not have to modify many jobs, you can make MIGRAT=YES the default on the ISPF install panel A.I.4.4.

If you have many data sets archived without MIGRAT=YES, you can recatalog them to MIGRAT using the RECATALOG SERVICE=CONVERT function of FDRARCH, as documented in [Section 51.61](#).

CONTINUED . . .

THIS PAGE WAS INTENTIONALLY LEFT BLANK

02.01 FDR DASD MANAGEMENT SYSTEM

**COMPONENTS
OF FDR**

The FDR DASD Management System consists of a number of components and options. The functions that you can perform with FDR programs depends on the terms of the license your organization has with Innovation Data Processing, the developer and vendor of FDR.

**BASIC
COMPONENTS**

Most licenses for FDR include these components:

- **FDR** (Fast Dump Restore) - performs full-volume backups and restores, as well as full-volume disk-to-disk copies. A utility (FDRTCOPY) is included to copy FDR backup tapes.
- **FDRDSF** (Data Set Functions)- performs backups and restores of selected data sets and VSAM clusters. It can also do data set restores from full-volume backup tapes produced by FDR.
- **FDRCOPY** - performs disk-to-disk copies of selected data sets and VSAM clusters.
- **SAR** (Stand-Alone Restore) - an IPLable program which does full-volume restores, so that you can restore FDR backups even when your MVS system is not working (such as restoring a failed disk which is required for MVS IPL). It can also do full-volume backups and can relabel disk volumes.

Note: if you are licensed **only** for FDRREORG, FDREPORT and/or FAST COMPAKTOR, you will not have the above components.

**OPTIONAL
COMPONENTS**

Your license may include some or all of the following optional components:

- **FAST COMPAKTOR** (FDRCPK) - reorganizes disk volumes to maximize free space for new allocations, and to reduce the number of extents occupied by data sets. COMPAKTOR usually takes only a few minutes per volume.
- **ABR** (Automatic Backup and Restore) - automates the management of DASD, providing regular full-volume and incremental (only changed data set) backups, space management by archiving unused or unneeded data sets with automatic recall if they are later required, and application backups to preserve all data sets required by a specific application. ABR includes a number of reporting and management programs.
- **FDRAPPL** (Application Backup) - FDRAPPL is a part of ABR, but it can also be licensed separately. It is used to backup and restore data sets by application.
- **FDRREORG** - automates the reorganization of PDSs (Partitioned Data Sets), VSAM clusters, and IAM (Innovation Access Method) files. Also adds support for PDS reorganization to FDRCOPY.
- **FDREPORT** - is a powerful and highly flexible reporting program, which can report on data from a number of sources (catalog, VTOC, VVDS, ABR information, DFSMSHsm BCDS/MCDS and more). The contents and format of the report can be specified by the user. Sorting and summaries are provided, and JCL or other control statements can be generated. ABR users are automatically licensed for FDREPORT, but it can be separately licensed.
- **FDRCLONE** - allows you to "clone" your production data onto a test system (such as a Y2K system or LPAR) or at a disaster/recovery site. Cloned data sets are restored only when a batch job or TSO user actually references them, so only required data is restored, which may be a fraction of your total data. FDRCLONE also includes a utility, FDRDRP, which reduces the elapsed time when restoring many volumes from ABR volume backups.

Note: it is possible to license FAST COMPAKTOR, FDRREORG, and/or FDREPORT separately, if that is the only function you need. However, ABR requires a FDR license with the basic components listed above. FDRCLONE requires that you are licensed for ABR.

02.01 CONTINUED

**OPTIONAL
FUNCTIONS**

Your license may include this optional function:

- **FDR InstantBackup** - enhances FDR and ABR to use features of certain DASD subsystems, such as the IBM RVA and StorageTek Iceberg/SVA Snapshot feature, the EMC Symmetrix Timefinder feature and the Hitachi Data Systems 7700E ShadowImage feature, to create point-in-time backups and copies of disk data. An image of the data on disk can be "frozen" at a particular point-in-time and backed up while updates to the original data continues.

The enhancements to FDR, FDRDSF, and FDRCOPY work with almost any disk subsystem and allow you to create offline point-in-time images of entire disk volumes. The point-in-time images can be backed up with FDR or FDRDSF or copied with FDRCOPY, even though they remain offline to MVS.

For Snapshot, Timefinder and ShadowImage, FDR InstantBackup also enhances FDRABR volume backups (full-volume and incremental backups) to provide instant backups. In addition, it enhances FDRCOPY and FASTCPK to use Snapshot and Timefinder to quickly copy data tracks. You must also be licensed for FDRABR or COMPAKTOR to get these benefits.

**STORAGETEK
HSDM AND
EXHPDM**

FDR includes support for 2 StorageTek performance-enhancing functions. Contact StorageTek or visit their web site (www.storagetek.com) for more information on these products.

- **HSDM** (High Speed Data Mover) is a hardware option available on StorageTek SVA disk subsystems. The SVA is a successor to the StorageTek Iceberg and IBM RVA disk subsystems and uses the same internal data structure. Data in these subsystems is internally compressed. A normal backup and restore requires that the control unit decompress the data during backup and recompress it during restore.

HSDM allows FDR to read and write the compressed track images directly, avoiding the overhead and reducing amount of data that must be sent over the disk and tape channels. HSDM reduces backup times by up to 60%, depending on your data. HSDM is invoked by the DCT=YES operand on any DUMP statement.

HSDM support requires that you be licensed for FDR InstantBackup.

- **ExHPDM** (High Performance Data Mover) is a software product from StorageTek which improves backup performance while reducing the number of tapes required. It does this by consolidating several concurrent backups onto one tape data set and formatting the data into larger blocks. It is designed for high-performance, high-capacity tape subsystems such as the StorageTek 9840, Timberline and Redwood, but it is hardware-independent and also supports IBM Magstar 3590s and other drives.

FDR V5.3 level 30 supports ExHPDM for FDR and FDRDSF backups. There are no special licensing requirements for ExHPDM with FDR and FDRDSF.

FDR V5.3 level 40 supports ExHPDM for FDRABR volume backups and application backups (FDRAPPL) but not for Archive Backups. ExHPDM support for FDRABR requires that you be licensed for FDR InstantBackup.

In either case, ExHPDM is invoked by a SUBSYS= operand on a TAPEX DD statement. Details are found in [Section 80.33](#) and the ExHPDM user documentation. Brief documentation is also found in each appropriate section (e.g., section 10 for FDR).

**IBM
SYSTEMPAC**

You may order FDR, including any of the optional features, as part of IBM's SYSTEMPAC. The SYSTEMPAC will include prebuilt libraries for FDR with the options you have chosen.

CONTINUED . . .

02.02 FDR AND FDRDSF INTRODUCTION

FDR Program FDR (Fast Dump Restore) provides full-volume backups and restores, as well as full-volume disk-to-disk copies. FDR is described in detail in [Section 10](#).

FULL VOLUME BACKUPS FDR backups operate on one DASD volume at a time, writing the full-volume backup to a tape or DASD data set; a separate backup data set is required for each DASD volume processed.

- The FDR execution JCL specifies the DASD volumes to backup, and the backup data set for each. The backup will contain an image of:
- the label track (Cylinder 0 Track 0), including IPL text for IPLable volumes
- the VTOC tracks
- the VTOCIX (Indexed VTOC) tracks
- the VVDS tracks if one is present (VSAM clusters and SMS data sets).
- all tracks which are allocated to other data sets and VSAM clusters, determined by analyzing the DSCBs (Data Set Control Blocks) in the VTOC. Tracks which are not currently allocated to any data set will not be backed up

an edited version of information from the VTOC and VVDS is also stored at the beginning of the backup, serving as an index of the data actually in the backup, and providing information for allocation of output data sets during data set restores.

FDR can optionally create two backup data sets concurrently, while reading the input disk volume only once. FDR can also optionally backup several disk volumes concurrently, using internal subtasking, as long as the backups are directed to separate devices.

Note: a common format is used for all backups in the FDR DASD Management System, including FDR, DSF, ABR and SAR (see notes below on data set backups); this means that FDR and FDRDSF can restore from tapes created by FDRABR, and vice versa. Actually, the FDR format has been modified several times since its initial release as support for new disk devices was added and performance improvements were made. Although newer formats may not be restorable by older releases of FDR, the current FDR can always restore older backups and is capable of producing backups which can be read by older releases.

WARNING: FDR format backups cannot be copied by standard copy utilities, such as IEBGENER. FDR includes a utility, FDRTCOPY, which is used to copy FDR backups.

FULL VOLUME RESTORES FDR restores a physical image of a disk from an FDR full-volume backup file. The DASD volume receiving the restore must be the same device type (e.g., 3390), but not necessarily the same size (number of cylinders) as the volume that was dumped. All the tracks in the backup data set will be restored to their original locations (cylinder and head address). Every restored track will look exactly like it did when dumped; data sets are not reorganized, reblocked or changed in any other way.

The output disk volume will usually become an exact image of the original disk including its volume serial, but you can optionally retain the volume serial of the output disk, and can rename key data sets (such as the VTOCIX and ABR Model DSCB) during the restore.

If the restore results in a volume with a volume serial which duplicates another online volume (e.g., when the original volume is still online), FDR will place the output volume offline. You must relabel one or the other volume before the next IPL, so a reminder message is sent to the operator's console.

WARNING: if the volume serial of the output volume is retained, any VSAM clusters on the volume will not be usable until the volume is renamed to its original serial. SMS-managed volumes are always renamed to the original volser from the backup when restored.

02.02 CONTINUED**DISK TO DISK
COPYING**

FDR allows you to copy one disk volume to another disk volume of the same type. Essentially this does a backup and restore all in the same job, so all considerations for restore apply. No backup data set is required, but you can optionally create a FDR backup during the copy operation.

FDRDSF

Program FDRDSF (Data Set Functions) performs backups and restores of selected data sets and VSAM clusters. It can also do data set restores from full-volume backup tapes produced by FDR, ABR and SAR, as well as from data set backups produced by ABR. FDRDSF is described in detail in [Section 20](#).

**DATA SET
BACKUPS**

FDRDSF backups are similar to FDR backups, using the same format, but containing only the information and data tracks associated with the data sets selected by DSF control statements. You may select data sets by specific name, by prefix, or by sophisticated data set name masking. DSF backups operate on one DASD volume at a time, reading the VTOC of the volume, backing up the data sets you have selected and writing the backup to a tape or DASD data set; a separate backup data set is required for every DASD volume processed but that backup data set may contain backups of many original DASD data sets. The backup will contain an image of:

- all tracks which are allocated to the selected data sets and VSAM clusters, according to the DSCBs in the VTOC.
- an edited version of information from the VTOC and VVDS is stored at the beginning of the backup, serving as an index of the data actually in the backup, and providing information for allocation of output data sets.
- The VTOCIX and VVDS may be backed up as data sets, if they are selected by DSF control statements. However, they are not required to be able to restore data sets from the backup, and they are usually not included in the backup.

Since FDRDSF does not support selection of data sets from your catalogs (CATDSN=) and requires that you know which volumes the desired data sets reside on, Innovation recommends that you use FDRAPPL (Application Backup) to backup data sets selected from the system catalogs.

**DATA SET
RESTORES**

FDRDSF restores are considerably different from FDR restores:

- the data tracks of the selected data sets may be restored to a different physical location (cylinder and head address) than the original tracks occupied. In other words, you do not have to be concerned about the location of the output data set or the number of extents it is in.
- data sets and clusters may be renamed during restore.
- if an output data set (either the original data set name or the new name) already exists, and is large enough, DSF will replace the contents of the existing data set and update its describing information (VTOC and VVDS information). By default, DSF will restore an existing data set to whatever volume it is currently cataloged.
- if an output data set does not exist, DSF will allocate it, automatically making it large enough to hold the contents of the input data set. Usually DSF will also catalog the output data set as well (if the output data set is already cataloged to another volume, DSF will not catalog it unless you instruct it to do so). VSAM clusters must be cataloged when they are allocated, so a VSAM allocation will fail if it is already cataloged (but you have an option to delete the existing cluster and catalog the new one in its place).
- DSF usually restores data sets to the same device type they were backed up from (e.g., from a 3390 backup to a 3390). This is called a like device or physical restore. The size (number of cylinders) on the original disk volume and the output volume is not important, as long as the output has enough free space to hold the output data sets being allocated. In a like restore, the original data tracks of the selected data sets are restored exactly as they were backed up (but there is an option to reblock certain data set types).

CONTINUED . . .

02.02 CONTINUED

- DSF can also restore data sets to a different device type (such as a backup of a 3380 to an output 3390), with some restrictions. This is called an unlike device or logical restore since the data records of the original data sets are usually reformatted to make better use of the track capacity of the output disk.
- Output data sets may be directed to various output disk volumes. Data sets from one backup, originally all on the same disk, can be restored to many output disks concurrently (reading the backup only once). FDRDSF JCL must point to the backup data sets, but the output volumes can be identified by JCL or by DSF control statements. You can also identify a list or group of volumes as the target volume; DSF will find one with sufficient space for the data set.
- Multi-volume data sets can be restored, but only to the same number of volumes they originally occupied when dumped. Multi-volume VSAM is handled, but only when restored to the original device type.
- At the end of the restore, DSF will update the DSCB of the output data set and, for VSAM and SMS-managed data sets, its VVDS entry, so that they properly describe the data that was restored.

FDRCOPY

Program FDRCOPY performs disk-to-disk copies of selected data sets and VSAM clusters. Essentially it does a FDRDSF backup and a restore all in the same step, without the requirement of creating a backup data set, so all the notes on FDRDSF restore above apply. FDRCOPY is described in [Section 21](#).

However, FDRCOPY has the ability to select the input data sets from the MVS catalog. You can specify a data set name, prefix, or data set name mask, and FDRCOPY will locate those data sets in the MVS catalog and automatically identify and process the DASD volumes required. You do not need to specify either the input or output disk volumes in the FDRCOPY JCL. You may skip the catalog search and search specified DASD volumes directly, if you wish.

If you are also licensed for FDRREORG, FDRCOPY includes a function for high-speed reorganization of PDSs (Partitioned Data Sets); this is also called PDS compression. The PDSs to be reorganized can be selected by normal FDRCOPY selection, from MVS catalogs or volume VTOCs.

**ABSOLUTE
TRACK
OPERATIONS**

FDRDSF and FDRCOPY can also dump, restore, and copy data tracks by their physical, absolute track address. Ranges of tracks to be processed are identified by a starting and ending track address (cylinder and head number). In this mode, the original tracks can be restored or copied only to the original track addresses; you cannot restore to a new location on the output disk with absolute track addresses.

**PRINTING
DISK DATA**

FDRDSF has a function to print the contents of disk data tracks. The printout shows the physical layout of each track, including each physical record and the hardware count field and key field (if present) for each record. You may select the data tracks to print by data set name and/or by absolute track.

CONTINUED . . .

02.02 CONTINUED

SAR SAR (Stand-Alone Restore) is an IPLable program which runs without an operating system, whose main function is to restore FDR full-volume backup tapes at a time when you cannot execute normal FDR because your MVS system cannot be IPLed. This could occur if disk volumes required for MVS IPL (such as residence volumes, paging volumes, spooling volumes or library volumes) are unavailable due to hardware or software problems. SAR can also be used to setup new data centers by restoring backups of volumes prepared at another site. SAR is described in detail in [Section 15](#).

There is an IPLable copy of SAR in the first file on every FDR distribution tape, but since that is a labeled tape, **five hardware IPL** functions are required to bypass the labels and load the program. Innovation recommends that you copy the SAR IPL text to an unlabeled tape or to one or more disk volumes for quick IPL when required. Once loaded, SAR can restore multiple FDR backup tapes without an additional IPL.

SAR can also create a full-volume backup of disk volumes on tape. This might be required to backup a non-critical volume so that a critical volume can be restored in its place. Although this backup is in normal FDR format (described earlier in this section), it does not contain formatted VVDS data, so DSF cannot restore VSAM or SMS-managed data sets from a SAR backup.

SAR can also relabel a disk volume.

FDR INSTANTBACKUP If you are also licensed for FDR InstantBackup, it enhances FDR, FDRDSF, and FDRCOPY to provide the ability to take backups which are frozen at a given point-in-time. This works with almost any disk subsystem. It allows you to capture a point-in-time image of an online disk to an offline disk, effectively preserving the image of the online disk at that point-in-time. FDR InstantBackup can then read the offline disk and create the required backups or copies, without relabeling the disk or bringing it online. FDR InstantBackup for various hardware platforms is described in [Sections 25-29](#).

CONTINUED . . .

02.03 COMPAKTOR INTRODUCTION

- FAST COMPAKTOR** FAST COMPAKTOR (FASTCPK) is an optional FDR component which quickly reorganizes DASD volumes, to maximize free space on a volume and reduce the number of extents in multi-extent data sets. Program FDRCPK is described in [Section 40](#).
- By default, FAST COMPAKTOR will analyze the current layout of the volume and will move data sets and data set extents on the volume to reduce multi-extent data sets and to gather free (unallocated) tracks into one or two large extents. It attempts to meet these objectives with minimum movement of data, so the COMPAKTion of a volume usually takes under 3 minutes.
- Although FAST COMPAKTOR is normally ordered as an option with the basic FDR components, it is possible to license COMPAKTOR by itself.
- COMPAKTOR** COMPAKTOR also has an option to take a FDR full-volume backup, and then COMPAKT the volume by restoring tracks from that backup. This is usually much more time-consuming than FAST COMPAKTOR, but it does include an option to move or expand the VTOC during the restore.
- MAPPING VOLUMES** As part of its normal operation, COMPAKTOR produces a map of the volume before and after the COMPAKTion, along with some statistics (number of free areas, size of free areas, etc.). The map may also be produced without doing an actual COMPAKTion. You can either map an online disk volume, or map a FDR backup tape.
- SIMULATION** All of the functions of COMPAKTOR can be simulated so that you can evaluate the need to do the real COMPAKTion and insure that the desired changes are being made. Simulation usually takes only a few seconds.
- FDR INSTANTBACKUP** If you are also licensed for FDR InstantBackup, and are executing FASTCPK on a volume in an IBM RVA or StorageTek Iceberg/SVA with the Snapshot feature or a EMC Symmetrix with the Timefinder feature, FASTCPK is enhanced to use hardware facilities of those systems to move data tracks during the COMPAKTion. This greatly reduces the elapsed time. FDR InstantBackup for the RVA/Iceberg/SVA and Symmetrix is described in [Sections 25](#) and [26](#).

02.04 FDRABR INTRODUCTION

ABR Program FDRABR (Automatic Backup Restore) automates the execution of FDR full-volume and data set backups and restores for the purposes of:

- **data availability** - creating backup copies of DASD data sets to protect against physical or logical loss.
- **space management** - identifying data sets which do not need to be on DASD (usually based on the last date they were used) and moving them to backups, from which they can be automatically recalled if needed. Data sets which will never be needed again can be scratched without creating a backup.
- **disaster recovery** - creating backups from which all or part of your DASD data can be quickly recreated at a disaster recovery site.
- This is accomplished by the use of several ABR functions, as described in the following:

ABR BACKUPS All the backup functions of ABR share these common characteristics:

- the backups are in standard full-volume (FDR) or data set (DSF) backup format, described in [Section 02.02](#). Data can be restored from ABR backups with FDR, FDRDSF, or ABR itself. SAR can be used to restore from ABR full-volume backups.
- although ABR backups still require a separate backup data set for each DASD volume processed, ABR will automatically stack multiple backup data sets on tape, creating multi-file tapes, to make best use of today's high-capacity tape volumes (such as IBM Magstar and StorageTek Redwood and 9840). If necessary, multiple output tape volumes are used. No special JCL is required since ABR will handle the file creation internally.
- ABR backups to disk are also supported. ABR will automatically allocate the backup data sets on the backup volumes you specify.
- the output devices (tape and/or disk) are specified in the ABR batch JCL. However, you only need to identify the output device; ABR will automatically name and create every backup data set and catalog it if required.
- You can specify in the JCL that two copies of each backup are to be created, even though the input disk is read only once. This is often used to create one copy for onsite recovery and a second for offsite storage for disaster recovery.
- If multiple output devices are specified in the JCL, ABR will automatically use internal subtasking to backup more than one DASD volume concurrently.

02.04 CONTINUED

**ABR VOLUME
BACKUPS AND
RESTORES**

ABR Volume Backups consist of full-volume and incremental backups which allow you to create daily backups of your disk volumes, so that you can do anything from restoring individual data sets to recreating entire volumes. However, you don't have to backup all the data every day. A daily full-volume backup of each of your DASD volumes would be too time-consuming, so users who do not have a product like ABR usually content themselves with weekend full-volume backups and leave themselves open to a weeks worth of updates lost if a failure or disaster occurs late in the week.

With ABR volume backups, you get the effect of daily full-volume backups without the elapsed time associated with them. You must still do periodic full-volume backups (most users do them on weekends, but you can schedule them to suit yourself; they don't even have to be weekly). On a daily basis (or sometimes more than once a day), you do an ABR incremental backup which only backs up the data sets which have changed or been created since the last backup of the volume, plus the label track, VTOC, VTOCIX and VVDS. These incremental backups are accumulated and are kept at least until the next full backup is executed.

If you have disk failures, or have to restore at a disaster site, ABR automates the restore process. It locates the most recent full-volume backup (or an earlier one if you prefer) and all the incremental backups that followed it. ABR then reads the most recently-created incremental, restoring the label track, VTOC, VTOCIX, VVDS and any data sets on that backup. It then reads back through the preceding incrementals and the full backup, restoring the most recent copy of each data set. The result is a volume which looks exactly like the original volume did at the time of the last incremental. All data sets are in their original locations, with the exact same allocation characteristics.

If you need to restore individual data sets, ABR tracks which full or incremental backup contains the most recent backup of each data set. Restore is as simple as providing the data set name; ABR will locate and dynamically allocate the proper backup and restore the data set; the restored data set can be renamed and/or directed to a new volume, but by default it is restored to its original volume with its original name. ABR has an option to track up to 13 previous backups for every data set, allowing you to easily restore earlier versions. ABR data set restores follow the same rules as DSF restores described in [Section 02.02](#).

ABR volume backups do not use a data base for recording the backups. Information about backups is stored in 4 places:

- every volume to be processed for ABR backups must have a special DSCB in the VTOC of the volume, known as the "ABR Model DSCB". It is called a model because the data set has no extents and occupies no tracks on the volume, similar to the model DSCB used with GDGs. The DSCB itself is used for recording data about volume backups for that volume, such as the date of the most recent backup. An ABR utility is used to create the ABR Model DSCB.
- the backup data sets created by ABR are cataloged in a standard MVS ICF catalog. They have a standard name format starting with a specified high-level index (default is FDRABR) which is usually assigned to a separate catalog, known as the ABR catalog.
- Information about the backups of individual data sets is stored in reserved fields of the Format 1 (F1) DSCB of each data set on disk.
- for data sets which have been scratched, so that their F1 DSCB is no longer available, ABR records backup information in a special MVS ICF catalog entry, called the "scratch catalog" entry, usually part of the ABR catalog.

During restore, ABR uses this information to construct a list of the backup data sets which must be read. In most cases, it will dynamically allocate the restore backup data sets (on tape or disk) and read them.

ABR volume backups are designed to allow you to recreate entire DASD volumes, or to restore backup copies of individual data sets. Details on ABR volume backups is found in [Section 50](#).

CONTINUED . . .

02.04 CONTINUED

**FDR
INSTANTBACKUP**

If you are also licensed for FDR InstantBackup, it enhances FDRABR Volume Backups to provide the ability to take backups which are frozen at a given point-in-time. This works with special facilities available on some disk subsystems, namely Snapshot on the IBM RVA and StorageTek Iceberg and SVA, Timefinder on the EMC Symmetrix, and ShadowImage on the Hitachi Data Systems 7700E. It allows you to capture a point-in-time image of an online disk to an offline disk, effectively capturing the image of the online disk at that point-in-time. ABR can then read the offline disk and create the required backups or copies, without relabeling the disk or bringing it online. ABR can also restore from the offline disk if necessary. FDR InstantBackup for Snapshot, Timefinder and ShadowImage is described in [Sections 25, 26 and 27](#).

Note that FDR InstantBackup may make it practical to do daily full-volume backups, since it moves the actual creation of the backup tapes outside of the backup window. This will use more tape than daily incrementals, but will greatly simplify restores and volume recovery.

FDRCLONE

If you are also licensed for FDRCLONE, you can use your regular ABR volume backups to recreate your production data on a test system (such as a Y2K test system) or at a disaster/recovery site, without restoring all of the data. On the test/recovery system, the program FDRCLONE marks the volumes or data sets you select as "cloned" in the system catalogs, but does not restore them. The first time that a batch job or TSO user references a cloned data set, it is restored from its most recent volume backup. So, only the data actually used at the test/recovery site is restored, which may be a fraction of your total data. The restores can be directed to new volumes, so you need to provide only the disk resources required to hold the data which is actually used. Details are found in [Section 50.70](#).

Your FDRCLONE license also includes FDRDRP (Disaster/Recovery Program), an ABR restore utility which manages full-volume recoveries from ABR volume backups (incremental and full-volume backups) by mounting each required backup tape as few times as possible, often only once per tape. This is usually used at disaster/recovery sites. Details are found in [Section 50.80](#).

**ABR
APPLICATION
BACKUPS AND
RESTORES**

ABR application backups (FDRAPPL), as their name implies, are designed to backup all the data sets that relate to a given application, on whatever disk volumes they reside. Since application data sets are usually cataloged, you can have ABR search the MVS catalogs for data sets matching one or more name masks, quickly and easily selecting the data sets and their cataloged volumes to be processed.

Application backup will process all of the data sets selected on the volumes selected. Since ABR always creates one backup data set per disk volume, it will automatically stack multiple backup files on the output tape, depending on how many disk volumes were read. Usually a single tape is used for output, creating one tape with all the selected data (a duplicate backup can be created).

Application backup does use a data base to record the data sets that were backed up. This data base is called an Application Control File (ACF); it has the same format as the Archive Control File described below. The ACF is very compact, recording several hundred data sets in a single track.

If data sets must be restored from an application backup, the information on their location is read from the ACF to build a list of backup data sets to be dynamically allocated and read. You can select the data sets to be restored, or simply restore the most recent copy of every data set in the ACF.

There are several ways of managing the ACF:

- you can create a new ACF every time that you execute application backup. It will contain only the data sets which were backed up in that ABR run. Usually a copy of the ACF will also be written to the backup tape, making the tape self-contained so that you can easily restore at a disaster site.
- you can use a permanent ACF for each application. This will record **all** backup copies of each data set which is part of the application. Usually you will need to backup and restore the ACF with another facility (such as ABR incremental backup) so that you can restore the application at a disaster site.

ABR Application backup provides for backup and restore of the data sets belonging to one application. [Section 52](#) describes this function in detail. You can license FDRAPPL separately without licensing all of ABR.

CONTINUED . . .

02.04 CONTINUED**ABR ARCHIVE
AND AUTO-
RECALL**

ABR Archive is used for space management, by moving data sets which are not currently required to be online from expensive DASD to less expensive medium such as tape or DASD (in a highly compressed format).

The DASD volumes to be processed by Archive must contain an ABR Model DSCB (described under "ABR Volume Backups" in this section); Archive does not record anything in the ABR model, but there is a flag enabling or disabling each volume for archiving. The data sets moved from DASD by ABR Archive are recorded in an Archive Control File (ACF). The ACF is a DSORG=DA (direct access) data set which can record several hundred such archived data sets in a single track. There is usually one common ACF for all archived data sets.

You can specify various criteria for selecting the data sets to be archived, such as "not cataloged" or "not used in the last 15 days". Different criteria can be specified for various data sets (by volume, by data set name mask, or a combination).

A data set which has been archived can be marked as eligible for automatic recall ("auto-recall"). With auto-recall, the catalog entry for an archived data set is left in the MVS catalog, with a special indicator that indicates that it was archived. When an archived data set is referenced by batch JCL, dynamic allocation, or TSO, an ABR catalog locate exit detects the indicator and automatically "recalls" the data set to disk before it is needed, so the use of archived data sets is transparent to the job or user.

TSO users normally have the option of bypassing the recall, and have a choice of waiting for the recall to complete (foreground) or doing the recall in a separate started task while they do other work (background). TSO users may also add recall requests (and some other ABR requests such as Archive) to a "remote queue" data set which is processed by a batch ABR job.

ABR Archive is often used to put one backup copy on disk and a second on tape. The disk backups will be in a highly compressed format, taking much less room than the original, yet they can be recalled quickly since no tape mount is required. Most installations expire the disk backup after a short time (15-30) days and either delete it or move it to tape.

ABR Archive can be used to implement the IBM concept of Tape Mount Management (TMM) where data sets which are normally on tape are directed to SMS-managed disk instead, and then moved to tape using ABR. This allows many such data sets to be stacked on tape and uses far fewer tape volumes.

ABR Archive is described in detail in [Section 51](#).

**ABR SUPER-
SCRATCH**

ABR Superscratch operates exactly like ABR Archive, but it does not backup the selected data sets, it simply deletes them. Superscratch can be used to delete data sets which you know will never be needed again (such as temporary data sets). There is a separate flag in the ABR Model DSCB to enable a volume for Superscratch.

ABR Archive and Superscratch is described in detail in [Section 51](#).

SIMULATION

Most functions of ABR can be run in simulation mode, allowing you to verify that the correct data will be selected when run for real.

ABR UTILITIES

ABR includes several utility programs, FDRABRM for setting options on disk volumes, FDRABRCM for maintaining the ABR catalog, FDRARCH for maintaining the Archive/Application Control File, FDRTSEL for automatic movement of ABR backups, and FDRABRP for simple reporting. ABR also includes FDREPORT (see below) for more sophisticated reporting.

CONTINUED . . .

02.05 FDREPORT INTRODUCTION

FDREPORT FDREPORT is a sophisticated, flexible and high-performance reporting package which can report on hundreds of characteristics of data sets or volumes in a format chosen by FDREPORT or specified by the user. FDREPORT is included with ABR, but it can also be licensed separately. FDREPORT is described in detail in [Section 54](#).

FDREPORT takes its input from a variety of sources:

- MVS catalogs - a data set name mask can be specified to quickly search catalogs for the required data sets
- VTOC of DASD volumes - information from the DSCBs of the data sets is extracted
- VVDS of DASD volumes - for VSAM clusters and SMS-managed data sets
- ABR backup information - for data sets processed by ABR volume backups
- Archive/Application Control File (ACF) - for archived data sets and those processed by application backups
- DFSMSHsm BCDS/MCDS data
- data recorded by previous FDREPORT executions

From the hundreds of different fields available, you can select a subset to be reported, and can optionally specify the positioning of those fields in the report, although FDREPORT does most of the formatting work for you. In addition to simple printed reports, FDREPORT can sort and summarize the data selected. It can also generate control statements and/or JCL for other utilities based on the data collected. Although most reports are based on individual data sets, you can report on the state of entire DASD volumes (such as volume free space).

So that you can immediately begin benefiting from the power of FDREPORT, 50 pre-constructed FDREPORT reports have been provided as part of the Innovation "Health Check". These enable you to report on the status of your entire DASD farm, on individual groups of data sets, and also to check for various error conditions. You will find them in the JCL library on the FDR distribution tape; member names begin with "HCHECK", such as HCHECK2.

02.06 FDRREORG INTRODUCTION

FDRREORG FDRREORG automates the reorganization of disk data sets which can benefit from this reorganization, such as VSAM KSDS (keyed) clusters, partitioned data sets (PDSs) and IAM files. IAM (Innovation Access Method) is a separately priced product from Innovation Data Processing; IAM is a transparent performance-oriented alternative for VSAM KSDS and ESDS clusters. FDRREORG is described in detail in [Section 30](#).

FDRREORG control statements identify the data sets and/or volumes to be processed, and also specify criteria for choosing the data sets to be reorganized. There is little benefit in reorganizing a data set which has had little update activity, so you can specify in various ways the "degree" of disorganization required to select a given data set.

For VSAM and IAM, FDRREORG creates a backup data set (using parameters you provide) and copies the data to the backup and back again to the original data set to reorganize it. If you save the backup data sets, FDRREORG can be used as a backup utility for VSAM and IAM; the advantage over a utility such as IDCAMS is that many such data sets can be backed up with simple control statements. The reorganized data sets are usually written back to the original volume, but you can override the output volume, using FDRREORG to move VSAM and IAM to new volumes while reorganizing them. Note that the FDRREORG backups are not in FDR format, they are simple sequential copies of the data in each data set.

For PDSs, FDRREORG invokes a function of FDRCOPY to compress the PDS in place. This function can also be executed directly under FDRCOPY if you are licensed for FDRREORG.

02.10 FDR FUNCTIONS AND FEATURES

This section documents many of the features and characteristics which are common to all or most FDR components. The term "FDR" in this section refers to all the FDR components.

**OPERATING
SYSTEM
SUPPORT**

FDR supports all versions of MVS, including:

OS/390 all releases

ESA all releases

MVS/XA all releases

They are referred to collectively as MVS in this manual. All functions described in this manual will work on all versions of MVS, except where a certain level of MVS is required for support of a particular device type or data set type. If new versions of OS/390 require changes in FDR, those changes will be available at or shortly after GA (General Availability) of the new OS/390 release.

FDR will run on a Sysplex, including a Parallel Sysplex, but does not use or require any Sysplex services.

**DISK DEVICE
SUPPORT**

FDR supports all IBM disk types supported by the MVS versions listed above, including disks which emulate one of the above, such as IBM RAMAC, EMC Symmetrix, and others.

Even though most currently marketed disk systems (from IBM and other vendors) use PC-style hard disks to store data internally and may use a RAID format for recovery after disk errors, they always emulate one or more of the standard disk types, usually 3390 or 3380.

FDR uses EXCP (Execute Channel Program) to read and write disks with its own CCWs and I/O techniques. It does not use any other access method for disk access (FDRREORG does use VSAM and IAM to read and write those types of data sets). During backups, FDR will usually read 1/2 of a cylinder, or a whole cylinder, in one I/O.

The original 3380s and 3390s came in several models with different capacities (number of cylinders). Many of the PC-disk/RAID systems allow you to define 3390s or 3380s which have a user-defined capacity different from those of the true disks. FDR supports all of these, and includes the required support for special requirements of some of these systems.

FDR full-volume restores or copies must always be done to a disk of the same type as the original (i.e., 3390 to 3390). You may restore/copy to a disk with the same capacity as the original, or a disk with a larger capacity. In the latter case, the VTOC will be updated to reflect the larger size, the DOS flag will be set in the VTOC to indicate that the volume free space is invalid, and a special allocation will be done to force DADSM to recalculate the free space (this allocation is designed to fail, so you can ignore allocation error message, IEC614I RC=16, in the joblog of the FDR restore job). If an indexed VTOC exists on such a volume, it will be disabled; you must execute the ICKDSF utility to rebuild the indexed VTOC of the volume after the restore (BUILDIX function).

NOTE: The data set that FDR allocates to rebuild the free space has a name that starts with "FDRABR.V". If your installation has a data security system, then the user running the restore job must be authorized to create this dsname.

Special procedures are required for an FDR restore to a smaller disk. However, COMPAKTOR can restore FDR backups to a larger disk or smaller disk as long as all of the allocated disk space fits on the output disk. COMPAKTOR will automatically update the volume free space and rebuild the Indexed VTOC.

FDR data set restores and copies can be done to either the same device type (e.g., 3390 to 3390, a like restore) or, in many cases, to a different device type (e.g., 3380 to 3390, an unlike restore). The device capacity does not matter as long as there is enough room for the data sets being restored. [See Section 80](#) for more information on moving data between different DASD devices.

02.10 CONTINUED

**TAPE DEVICE
SUPPORT**

FDR supports all current tape types supported by the MVS versions listed above, including:

IBM 3590 (Magstar) high-capacity cartridge drives

IBM 3490E cartridge drives

IBM 3480/3490 cartridge drives

IBM 3420/3422 round reel drives

StorageTek Redwood and Timberline cartridge drives

StorageTek 9840 cartridge drives

Other drives which emulate one of the above

Automated tape libraries from various vendors, including the IBM 3494/3495 tape libraries and StorageTek Nearline silos, are supported. FDR also supports the IBM Virtual Tape Server (VTS), StorageTek Virtual Storage Manager (VSM) and similar hardware and software virtual tape systems from other vendors.

FDR may use BSAM (Basic Sequential Access Method) or EXCP (Execute Channel Program) to read and write tapes (as well as backups on sequential disk data sets).

All cartridge drives (except some early 3480 drives) support hardware compression, originally called IDRC (Improved Data Recording Capability). FDR supports tape hardware compression. Compression must be requested on tape DD statements in FDR jobs (TRTCH=COMP) unless it is the default in your installation. Since it increases the effective capacity of a cartridge, it is recommended.

CONTINUED . . .

02.10 CONTINUED**ERROR
PROCESSING**

During backup and restore operations, FDR continually analyzes the processing and reports any unusual conditions detected. Physical and logical errors are handled, including:

- physical I/O errors reported by the disk hardware, such as data checks. FDR prints out details on the error and on the I/O being executed. During backup, tracks which get I/O errors are not written to the backup; during restore, the results are unpredictable. Depending on options, FDR may accept a number of such errors before aborting the operation
- physical I/O errors on the backup data set (disk or tape). FDR prints details of the error. For tape, it will attempt to close the current tape and write the failed data to a new tape volume; depending on the nature of the tape error, this may not be possible.
- invalid disk track format, including non-standard or invalid Record 0 (R0) format. This usually indicates a data track that does not meet IBM standards and cannot contain useful data so the track will not be written to the backup.
- incorrect count fields where the CCHH (cylinder and head number) in the count field of a record does not match the physical CCHH of the data track. Since this is not necessarily an error (VM-formatted volumes containing mini-disks may have this condition), the number of such tracks is simply counted.
- I/O errors in the VTOC, including physical errors and logical errors such as the wrong number of DSCBs on a track. VTOC errors always terminate data set backups; for full-volume backups it forces FDR to backup all physical tracks on the volume.
- I/O errors in the VVDS. The backup will continue but VSAM and SMS-managed data sets may not be restorable.

On modern DASD volumes (especially RAID subsystems) and modern tape drives, physical I/O errors are rare. Most error recovery is done in the disk or tape control unit, and recovered errors are never even presented to FDR.

VSE VOLUMES

VSE (DOS) volumes are similar to MVS volumes, but may not follow all of the same rules. Any VSE volume which can be mounted to MVS system can be dumped or restored with FDR or FDRDSF. However, ABR functions should not be used.

CONTINUED . . .

02.10 CONTINUED

VM VOLUMES FDR can dump and restore volumes belonging to a VM system, both from MVS systems running as VM guests, and from independent MVS systems which can access the disk drives in use by VM. There are several variations which you must be aware of:

- if the volume was formatted by VM (such as VM residence, spool or paging volumes or volumes containing CMS mini-disks), the volume will contain a dummy MVS VTOC on CYL 0 TRK 0 so that it can be mounted on MVS systems. Since a MVS-formatted VTOC will never start at that location, FDR recognizes that dummy VTOC and automatically dumps or restores all of the physical tracks on the disk. Since there are no data sets in that VTOC, DSF can only do ABSOLUTE TRACK ADDRESS operations on such volumes. An individual CMS formatted mini-disk may be restored from an FDR or DSF full-volume backup if you know its extents on the volume, but only to its original disk locations; DSF cannot be used to move CMS mini-disks.
 - COMPAKTOR and ABR operations cannot be done on such disks.
 - Other FDR functions can be done by an independent MVS system.
 - FDR functions cannot be done by a MVS guest under VM since it cannot access the full VM volume.
- if the volume is defined as several mini-disks (with the first mini-disk starting at physical cylinder 0) and each is formatted as a MVS volume
 - all FDR functions can be executed against each mini-disk by a MVS guest running under VM. The MVS virtual machine must have each mini-disk defined as a device.
 - from an independent MVS system, MVS will see only the first mini-disk on the physical volume (it is possible to use DSF physical track processing to include the additional mini-disks, but this is not recommended).
- if the volume is defined as MVS mini-disks, but the first mini-disk starts on physical cylinder 1 or later,
 - an independent MVS system will treat it as a VM-formatted volume and backup/restore all physical tracks
 - for a MVS guest under VM, each mini-disk should be defined to the MVS system so it can be processed normally.
- If the volume is defined as a full-volume MVS-formatted disk, it can be processed normally from either a MVS guest (where the disk device must be dedicated/attached or defined as a full-volume mini-disk) or from an independent MVS system.

Stand Alone Restore (SAR) has the ability to do backups and restores on any VM virtual machine without an operating system. You must IPL SAR on a virtual machine just like you IPL CMS or MVS. SAR can process individual mini-disks (OS or CMS formatted) if they are defined to the virtual machine, and can be used to move a mini-disk to another of the same size (SAR absolute track operations are used for CMS mini-disks). [See Section 15](#) for details.

CONTINUED . . .

02.10 CONTINUED

**CHANGING
THE VOLSER
OR VTOC
DURING
RESTORE**

If a full-volume RESTORE or COPY changes the volume serial of the output volume or changes the location of the VTOC or Indexed VTOC, FDR will automatically update the proper MVS control blocks and pointers with the new information, on the MVS system it was executed on. If the new volume serial duplicates the serial of another online volume, the output disk will be placed offline and a warning message issued to the MVS operator.

WARNING: If a full-volume RESTORE or COPY is done on a disk shared with another MVS system (shared DASD), it is safest to VARY the volume offline to all other systems before the restore/copy. After the restore is completed, the MOUNT command can be issued on other systems to make it available again. If FDR changes the volser or VTOC/VTOCIX location on a shared disk, a warning will be issued to the operator since the volume will not be usable on other systems until this is done.

**CACHE
SUPPORT**

If an FDR backup is executed for a disk connected to a control unit with cache, FDR will use a technique which will minimize the impact of the backup on the cache. Normally, all the tracks read by an application are retained in cache. Since FDR reads many tracks and does not need to read them again, this may displace other tracks needed by other applications, impacting their performance. FDR instead uses a technique called "inhibit cache loading" which will not retain the tracks read in cache; however, if required tracks are already in cache they will be read from cache and left there. So, tracks currently in the cache belonging to this volume and other volumes will not be disturbed.

Some disk subsystems always read data into cache before sending it to the channel. On those, "Inhibit cache loading" usually means that the track is discarded from cache as soon as it is read, which provides the same result.

For restores, normal cache management occurs. If volumes are enabled for DASD FAST WRITE, tracks are written by FDR go into cache and NVS (Non-Volatile Storage) and will be written to the physical disk when the control unit has an opportunity. For small restores, DASD FAST WRITE will speed up the restore, since it doesn't have to wait for the physical disk. For large restores, the capacity of the cache/NVS is usually exceeded, so the performance benefit is minimal. However, for certain disk types (such as some IBM RAMAC models) DASD FAST WRITE is used by the control unit to circumvent some write performance limitations on RAID disks (the "RAID Write Penalty").

**SUPPORTED
DATA SET
TYPES**

FDR supports backup and restore of all types of data sets which can reside on disk on supported MVS systems, including:

- Physical Sequential (PS)
- Partitioned (PDS)
- Direct Access (DA)
- ICF VSAM (ESDS, KSDS, RRDS, VRRDS, Linear, Alternate Indexes)
- DB2 data files and catalogs
- IAM (Innovation Access Method) – an alternate to VSAM KSDS and ESDS
- PDSE – an alternate to PDSs
- Physical Sequential Extended (PSE) – extended format PS data sets on SMS volumes (may be striped and/or compressed)
- Extended Format KSDS VSAM – on SMS volumes, may be compressed and may exceed 4GB in size
- Hierarchical File System (HFS) – used with UNIX System Services (Open Edition). FDR can backup/restore entire HFS data sets but not individual files within the HFS data set.

See [Section 80](#) for details on the FDR support for these various data set types.

OAM, the Object Access Method, stores collections of data called "objects" (such as scanned images). OAM files may exist on DASD or on optical storage disks. OAM files on DASD are DB2 files, which are supported by FDR. OAM files on optical disks do not look like traditional DASD, and are accessible only through OAM facilities. Since optical files are WORM (Write-Once, Read-Many) they are never updated; OAM provides for creating a one-time backup of optical files on other optical disks.

CONTINUED . . .

02.10 CONTINUED

SMS SUPPORT FDR supports backup and restore of volumes and data sets that are under the control of System Managed Storage (SMS). This includes those data sets which *must* be SMS-managed, such as extended format sequential data sets and KSDS clusters. See [Section 70](#) for the special considerations in using FDR, DSF and ABR with SMS volumes.

FDR PERFORMANCE FDR backups and restores are primarily I/O bound, meaning that they use very little CPU time unless you use FDR software compression. Performance of FDR operations is limited primarily by the I/O capacity of the tapes and disk and of the channel paths connecting them to the CPU; performance may also be affected if multiple FDR operations or other applications are using the same channel paths or control units. Many factors affect performance, so it is difficult to make predictions about the performance you can achieve.

When trying to improve performance of FDR or any I/O bound application, tuning is usually a process of moving from one bottleneck to another. As you make changes to reduce one limitation, another becomes the new limiting factor.

CHANNELS: your disks and tapes may be connected by parallel channels (also called bus and tag channels) which transfer data at a maximum rate of 4.5MB/sec, by serial ESCON channels which transfer at up to 17MB/sec (3490E drives on ESCON channels are limited to 9MB/sec) or by serial fiber FICON channels at 100MB/sec. Obviously, if your tape and disk control units support ESCON or FICON, these channels are desirable. ESCON/FICON speeds are usually faster than that of the disks and tapes themselves, so the channel is rarely a bottleneck. However, your channel configuration, especially if it includes channel directors, may limit concurrent transfers.

CONTROL UNITS: a control unit is required to access an I/O device; some control units are separate (e.g., IBM 3990) and some are integrated (e.g., EMC Symmetrix). A control unit can control from 1 to 256 devices; the fewer devices behind a given control unit, the better performance will be when doing concurrent operations.

The control unit may connect to a CPU over multiple channels. Most disk control units can do concurrent data transfers over each channel connection allowing you to concurrently access as many disks as you have channels to the control unit. Most tape control units can do only one channel data transfer at a time, despite multiple channel connections, which limits the ability to do concurrent operations. Many tape subsystems have multiple control units allowing several concurrent transfers. Some StorageTek cartridge systems have a control unit per tape drive to improve performance.

DISK PERFORMANCE: When data required by FDR must be read from the physical disk (not in cache), performance is limited by the speed of that disk. For example, physical 3390 disks transfer at a maximum of 4.2MB/sec. However, transfer rates from systems which emulate 3380s and 3390s (e.g., IBM RAMAC and EMC Symmetrix) may be much higher. When required data is already in cache, it is usually transferred at channel speeds (e.g., 17MB/sec for ESCON). If your disk is capable of HSDM, you can reduce the data transfer time up to 60% ([see Section 80.33](#)).

TAPE PERFORMANCE: tape transfer rates are limited by the physical tape drives. IBM 3480s and 3490Es transfer at 3.0MB/sec; hardware compression (IDRC) typically gets 50% compression, which increases the effective rate to 6MB/sec. However, only one 3480/3490E drive can be actively transferring data for each tape control unit. IBM 3590 Magstars can transfer at 9MB/sec, and their hardware compression typically gets 70% which makes their effective rate about 22MB/sec (which may really be 17MB/sec due to ESCON channel limitations). IBM 3490E and StorageTek 9840 drives are even faster.

CONTINUED . . .

02.10 CONTINUED

**FDR PERFORMANCE
(continued)**

OUR RECOMMENDATIONS: it is difficult to predict what sort of performance you will get, but here are some guidelines. Most customers want to complete their backups in the minimum possible time; this is usually accomplished by running multiple backups concurrently (either as multiple jobs, or using the internal subtasking options of FDR and ABR).

If possible, concurrently backup disks which are on separate disk control units, or at least don't backup more disks from a given control unit than there are channel connections.

If possible, backup to tape drives which are on separate tape control units, on separate channels. If necessary, you can direct up to 3 backups through the same control unit (with some degradation) but more than 3 will seriously degrade all the backups. Use tape hardware compression (e.g., IDRC) if available. You may want to use FDR software compression (COMPRESS= operand) on parallel channels (to improve the effective rate of the channels) but **do not** use software compression on ESCON or FICON channels.

On 3480/3490 tape drives, you may want to allocate 2 output drives to each backup, using a JCL parameter such as UNIT=(3480,2) to eliminate the time required to rewind and unload the last tape and mount a new one (this can take over 1 minute). On 3490E and 3590 drives, rewind time is not a problem since the tape is already back at its load point when full.

If available, Automatic Cartridge Loaders (ACLs) should be set to "system" mode and loaded with scratch tapes so that no operator intervention is required between tapes.

There are no special considerations for Automated Tape Libraries (ATLs, also called "silos"). However, if the tape drives are 3480, 3490 or 3490E drives, and sufficient drives are available, you may want to specify UNIT=(type,2) to allocate 2 output tape drives during backups, to avoid delays at EOJ while the robot changes tapes.

**TAPE VOLUME
USAGE**

This table shows the maximum number of tape volumes required to hold a full-volume backup of disks of various capacities, assuming that the disks are full. These calculations are based on the compressed capacity of each tape type assuming an average hardware compression of 50%. Your results will vary depending on the compression results of your data and the amount of data to be backed up.

<u>Disk type and model</u>	<u>3480 volumes (400MB)</u>	<u>3490E standard volumes (800MB)</u>	<u>3490E enhanced volumes (double length, 1.6GB)</u>	<u>StorageTek EETape (quadruple length, 3.2GB)</u>
3390-9 (8.5GB)	22	11	5	3
3390-3 (2.8GB)	7	4	2	1
3380-K (1.8 GB)	5	3	2	1

02.10 CONTINUED

**TAPE VOLUME
USAGE**

High-capacity cartridges can hold the backup of many disk volumes. This table shows the approximate number of backups which will fit on these cartridges, assuming an average hardware compression of 70%.

<u>Disk type and model</u>	<u>IBM Magstar 3590 standard volume (30 GB)</u>	<u>IBM Magstar 3590E standard volume (60 GB)</u>	<u>IBM Magstar 3590E enhanced volume (120 GB)</u>	<u>StorageTek 9840 (80 GB)</u>	<u>StorageTek Redwood (30 GB)*</u>
3390-9 (8.5GB)	3	7	14	9	3*
3390-3 (2.8 GB)	10	21	42	28	10*

* based on a 10GB (uncompressed) Redwood cartridge. StorageTek Redwood also supports cartridges with uncompressed capacities of 30GB and 50GB, so multiply the numbers above by 3 or 5 if they are used.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

10.01 FDR OVERVIEW

FDR Program FDR (Fast Dump Restore) provides full-volume backups and restores, as well as full-volume disk-to-disk copies.

FULL VOLUME BACKUPS FDR backups operate on one DASD volume at a time, writing the full-volume backup to a tape or DASD data set; a separate backup data set is required for each DASD volume processed.

The FDR execution JCL specifies the DASD volumes to backup, and the backup data set for each. The backup will contain an image of:

- the label track (Cylinder 0 Track 0), including IPL text for IPLable volumes
- the VTOC tracks
- the VTOCIX (Indexed VTOC) tracks
- the VVDS tracks if one is present (VSAM clusters and SMS data sets).
- all tracks which are allocated to other data sets and VSAM clusters, determined by analyzing the DSCBs (Data Set Control Blocks) in the VTOC. Tracks that are not currently allocated to any data set will not be backed up
- an edited version of information from the VTOC and VVDS is also stored at the beginning of the backup, serving as an index of the data actually in the backup, and providing information for allocation of output data sets during data set restores.
- some information contained only the catalog entry for the data sets being backed up will also be included in the backup. This currently includes PATH definitions for the alternate indexes (AIXs) of ICF VSAM clusters and the aliases of ICF VSAM user catalogs. Other catalog-only information, such as OWNERID, is **not** backed up and so will not be recreated during a data set restore.

FDR can optionally create two backup data sets concurrently, while reading the input disk volume only once. FDR can also optionally backup several disk volumes concurrently, using internal subtasking, as long as the backups are directed to separate devices.

Note: a common format is used for all backups in the FDR DASD Management System, including FDR, DSF, ABR and SAR (see notes below on data set backups); this means that FDR and FDRDSF can restore from tapes created by FDRABR, and vice versa. Actually, the FDR format has been modified several times since its initial release as support for new disk devices was added and performance improvements were made. Although newer formats may not be restorable by older releases of FDR, the current FDR can always restore older backups and is capable of producing backups which can be read by older releases.

WARNING: FDR format backups cannot be copied by standard copy utilities, such as IEBGENER. FDR includes a utility, FDRTCOPY, which is used to copy FDR backups. FDRTCOPY is documented in [Section 60](#).

10.01 CONTINUED

**FULL VOLUME
RESTORES**

FDR restores a physical image of a disk from an FDR full-volume backup file. The DASD volume receiving the restore must be the same device type (e.g., 3390), but not necessarily the same size (number of cylinders) as the volume that was dumped. All the tracks in the backup data set will be restored to their original locations (cylinder and head address). Every restored track will look exactly like it did when dumped; data sets are not reorganized, reblocked or changed in any other way.

The output disk volume will usually become an exact image of the original disk including its volume serial, but you can optionally retain the volume serial of the output disk, and can rename key data sets (such as the VTOCIX and ABR Model DSCB) during the restore.

If the restore results in a volume with a volume serial which duplicates another online volume (e.g., when the original volume is still online), FDR will place the output volume offline. You must relabel one or the other volume before the next IPL, so a reminder message is sent to the operator's console.

WARNING: if the volume serial of the output volume is retained, any VSAM clusters on the volume will not be usable until the volume is renamed to its original serial. SMS-managed volumes are always renamed to the original volser from the backup when restored.

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

**DISK TO DISK
COPYING**

FDR allows you to copy one disk volume to another disk volume of the same type. Essentially this does a backup and restore all in the same job, so all considerations for restore apply. No backup data set is required, but you can optionally create a FDR backup during the copy operation.

WARNING: before you do a full-volume copy you should make sure that the target volume is offline to all systems other than the system where the copy is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

**FDR
INSTANTBACKUP**

If you are also licensed for FDR InstantBackup, it enhances FDR to provide the ability to take backups which are frozen at a given point-in-time. This works with almost any disk subsystem. It allows you to capture a point-in-time image of an online disk to an offline disk, effectively preserving the image of the online disk at that point-in-time. FDR InstantBackup can then read the offline disk and create the required backups or copies, without relabeling the disk or bringing it online. FDR InstantBackup for various hardware platforms is described in [Sections 25-29](#).

FDR InstantBackup also enables support for the HSDM (High Speed Data Mover) feature available on some disk subsystems. HSDM allows FDR to backup and restore data in an internal compressed format. If your disk subsystem includes HSDM, you can invoke HSDM support by adding the DCT=YES operand on DUMP statements. [See Section 80.33](#) for details.

10.02 FDR PROCESSING OPTIONS AND REQUIREMENTS

**DISK BACKUP
OPTIONS**

FDR has three backup processing modes in a single FDR jobstep: dump one disk volume, dump up to thirty-nine disk volumes serially, or dump up to thirty-nine disk volumes to up to nine tape drives concurrently.

A. SINGLE VOLUME MODE

Dump the single disk volume referred to by the DISKx DD statement to the tape referred to by the TAPEx DD statement. This is the default mode of FDR when a single DISKx/TAPEx pair is provided in the JCL.

B. SERIAL MODE

Sequentially (one at a time) dump all the disk volumes referred to by DISKx DD statements to the tapes referred to by the corresponding TAPEx DD statements; backups are done in the order that the DISKx DDs appear in the JCL. Serial Mode is also the default when multiple DISKx/TAPEx pairs are provided.

C. ATTACH MODE

Concurrently dump the disk volumes referred to by the DISKx DD statements to the tapes referred to by the corresponding TAPEx DD statements. To invoke Attach Mode, you must specify the ATTACH or MAXTASKS= operand on the DUMP statement, or PARM=A on the EXEC JCL statement.

Although there may be up to thirty-nine TAPEx/DISKx pairs, the TAPEx DDs must not specify more than nine unique tape units. The TAPEx DDs may include UNIT=AFF or VOL=REF to use the same drive for multiple backups or to place multiple backup files on the same tape volume; the number of unique units determines the number of concurrent DUMPs unless limited by MAXTASKS=.

FDR will serialize the backup of any disk volumes whose TAPEx DDs point to the same tape unit.

**DISK RESTORE
OPTIONS**

When FDR is restoring a backup to a volume with the same volume serial as the original, e.g., restoring volume ABC123 to a disk with volser ABC123, the volume is simply restored and the contents of the volume are replaced with data from the backup. This is often the case when restoring a volume back to its previous contents, overlaying the current contents of the volume.

But when the output volume for a restore currently has a different volume serial from the serial of the disk in the FDR backup data set, FDR has two options:

A. RESTORE THE INPUT VOLUME SERIAL TO THE RECEIVING VOLUME

To change the Volume Serial Number of the output disk to match the serial of the disk that is on the backup tape, use the value PARM=R on the EXEC JCL statement, or specify CPYVOLID=YES on the RESTORE command.

B. RESTORE RETAINING THE VOLUME LABEL OF THE RECEIVING VOLUME

To retain the Volume Serial Number of the output disk, use the value PARM=N on the EXEC JCL statement, or specify CPYVOLID=NO on the RESTORE command or use a RESTORE command with the CPYVOLID operand omitted. If you intend to eventually relabel this volume back to the original serial number of the restored volume, you should also specify VOLRESET=NO, but if you intend to use the restore data on the output volume's serial, specify VOLRESET=YES.

If multiple restores are requested (multiple TAPEx/DISKx pairs), the restores are done serially in the order that the TAPEx DDs appear in the JCL unless MAXTASKS= is specified to request multiple concurrent restores. Similar to backup, FDR will recognize if certain TAPEx DD statements point to the same tape unit and will serialize restores from that unit.

10.02 CONTINUED

COPY OPTION

FDR can copy up to thirty-nine disk volumes to new disk volumes directly without creating a backup, in one FDR jobstep. FDR will copy the volume specified by the DISKx DD statement to the disk volume specified by the matching TAPEx DD statement. As with a FDR RESTORE the label of the volume being copied to may be retained or changed. The duplicate TAPExx DD statement is supported and will result in a backup being taken of the volume in addition to the copy.

MEMORY REQUIREMENTS

For backups specifying RTC=YES or DCT=YES, FDR requires:

- About 500K of storage below the 16MB line for programs and control blocks.
- About 2MB (2048K) of above-the-line storage for each concurrent backup.
- If COMPRESS= is specified, about 100K of below-the-line storage for each concurrent backup.

For backups specifying neither RTC=YES nor DCT=YES, FDR requires

- About 500K of storage below the 16MB line for programs and control blocks.
- About 1MB (1024K) of below-the-16MB-line storage for each concurrent backup.
- If COMPRESS= is specified, an additional 1M of below-the-line storage for each concurrent backup (total 2MB per backup).

In single volume and serial backup mode, FDR executes only one backup task. If ATTACH is specified, FDR will do up to 9 concurrent backups, limited by the number of TAPEx/DISKx DD statement pairs in the step JCL.

For example:

1 backup (or serial backups) without RTC/DCT – 1500K below the line

3 concurrent backups without RTC/DCT – 3500K below the line

3 concurrent backups with RTC or DCT – 500K below the line, 6MB above the line.

An FDR full-volume restore requires about 300MB of above-the-line storage for each concurrent restore.

FDR always uses the exact memory it requires for a given function, no matter what REGION= value is specified. So, if the region is too small, FDR will fail; if it is too large, FDR will not use the excess so a large region has no negative impact. For this reason, you may want to specify REGION=0M in FDR JCL to request the largest available region. Alternately, you can specify REGION=xxM (xx greater than 16) to get the largest available region and the specified amount of above-the-line region; the default above-the-line region is 32M unless your installation has changed it.

COMPRESS OPTION

FDR can be instructed to compress the data on the sequential backup file using Innovation's own proprietary software compression algorithm.

It is **not recommended** for backups to tape attached by ESCON or FICON channels because of the speed of the channel. Software compression will be ignored for backups created with the HSDM disk hardware option (DCT=YES) since the data is pre-compressed.

NOTE: FDR restore, including SAR, will automatically recognize a compressed backup file during a restore operation. No special option is required to restore a compressed backup.

DUPLICATE TAPE OPTION

FDR has an option to create a duplicate or second copy of the backup tape during dump processing. When several volumes are dumped duplicate backup files may be made for one or more of the disks regardless of the others.

While dumping a disk to a TAPEx DD statement, the duplicate backup will be written to the TAPExx DD statement (same "x" value twice) if it is present. You may have TAPExx DDs for some TAPEx DDs and not for others in the same step.

Memory requirements do not increase with the use of the duplicate tape option.

10.02 CONTINUED

STORAGETEK FDR supports the ExHPDM (High Performance Data Mover) software product from StorageTek.
EXHPDM ExHPDM takes multiple concurrent tape outputs (such as FDR backup TAPEX or TAPEXX DD
SUPPORT statements) and directs them to a smaller number of tape drives, interleaving the data in a single tape file. ExHPDM is invoked by adding the SUBSYS= operand to the TAPE DD statements. See [Section 80.33](#) for more details.

SECURITY Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system that supports the SAF (Security Authorization Facility) interface such as ACF2 or TOP SECRET, you can enable the ALLCALL option. For FDR this results in these security checks:

- for DUMP, FDR will check to see if your userid has at least READ authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority or the input volume is not protected by your security system, then FDR will check if you have at least READ authority under the DATASET security class to every data set being backed up. If you don't have authority to the entire volume or authority to every data set on the volume, the backup is terminated.
- for RESTORE, FDR will check to see if your userid has UPDATE authority to the entire output volume under the DASDVOL class. If not, the restore is terminated. But if the output volume is not protected by your security system, the restore will be done with no additional security checks. For this reason, Innovation recommends defining volume-level security rules to control full-volume restores.
- for COPY, FDR will do the DUMP checks above on input volumes and the RESTORE checks above on output volumes.

CONTINUED . . .

10.02 CONTINUED

**DATA SET
ENQUEUE
OPTION**

You can request, via the DSNENQ= operand, that each data set being dumped or copied be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name of SYSDSN and a minor name of the data set name itself, for each data set found in the VTOC of the input disk; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. Note that FDR cannot tell if the data set is being used for input or output by the other task. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations.

If you have requested data set ENQs, any data set that is in use will cause a FDR158 warning message to be printed; this will set the job error flag and will cause a U0888 abend when the step is complete (see "Step Termination" below). If you don't want in-use data sets to be considered an error, specify the ENQERR=NO operand; this prints the FDR158 message without setting the error flag.

Optionally you can request that inactive data sets be enqueued to FDR during the backup, to insure that no other job or TSO user can access the data set until the backup is done.

In-use data sets will still be dumped or copied, since this is a full-volume operation, but you must be aware that the backups of data sets which are being updated during the backup may be unusable, depending on the nature and format of the data.

The DSNENQ= operand has 4 possible values:

USE -- data sets will be enqueued for the duration of the backup or copy from this disk volume. For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option.

TEST -- data sets will only be tested to see if they are enqueued to another task at the time that the backup or copy from this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the dump is proceeding.

HAVE -- The data sets will be enqueued for the duration of the dump. If a data set is in use, the MVS operator must interact with FDR to decide how to proceed; a message (FDRW27) is issued to the MVS console, and the operator can respond:

WAIT - wait for the data set to become available; if it is not eventually dequeued, the FDR job may time out, so the operator must not reply WAIT for data sets in use by long-running jobs or tasks such as transaction processing systems like CICS.

NOWAIT - do not enqueue the data set. The FDR158 warning message is issued.

RETRY - try the enqueue again. If it fails again, the FDRW27 message is reissued.

NONE -- No data set ENQ will be issued. This is the default.

NOTE: If a data set name appears in a DD statement with DISP=SHR within the FDR job (not necessarily in the FDR step, FDR will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the FDR job ends.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

There is no option to ENQ data sets during a full-volume restore.

Use DSNENQ= to prevent other tasks from updating (or reading) data sets being dumped. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

If HFS=QUIESCE is specified, special processing is done for HFS data sets (Hierarchical File System, used by OS/390 Unix System Services). If the SYSDSN ENQ cannot be acquired, this may mean that the file system is mounted to OS/390 Unix, so FDR will attempt to quiesce the file system during the backup. Details on the quiesce function are found in [Section 80.11](#).

10.02 CONTINUED

**VTOC
ENQUEUE
OPTION**

FDR also supports, via the ENQ= operand, an ENQ on the VTOC of every volume being dumped, restored, or copied. For shared DASD, it can also invoke a hardware RESERVE on the volume during the FDR operation.

The VTOC is protected by an ENQ with major name SYSVTOC and a minor name of the volume serial. This ENQ is held by any task doing updates to the VTOC, including allocation of new data sets, extension of data sets to new extents, and scratching of existing data sets. This ENQ is normally short duration, just for the few seconds necessary to update the VTOC, so if the ENQ is currently held by another task, FDR will wait for it to be released.

The SYSVTOC ENQ does **not** prevent access to existing data sets on the volume; it only insures that no other task is updating the VTOC while FDR is processing it. VTOC changes during a backup or copy could result in an invalid backup.

For disks shared with another MVS system or LPAR, ENQ=RESERVE requests that, in addition to the ENQ described above, a hardware RESERVE is done on the volume. RESERVE will prevent any system from reading or writing data on the volume, except for the system that FDR is running on, where only the ENQ protection applies. If you have a cross-CPU ENQ facility, such as GRS or MIM, you may be able to convert the RESERVE into a cross-CPU SYSVTOC ENQ and allow access to the volume during the operation (lookup SYSVTOC in the documentation for your product).

Use ENQ= to prevent other tasks from making changes to the VTOC during the backup or restore. For FDR backups and copies, since they do not modify the input disk volume, DSNENQ= (above) is usually more important than ENQ=. Member ENQ in the FDR ICL (Installation Control Library) has more information on VTOC ENQs.

**STEP
TERMINATION**

If no errors occur during the execution of FDR, the FDR jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

- Some errors are critical. The jobstep ends immediately with a user ABEND.
- Some errors are critical only to a particular operation. For example, during a backup, some errors cause the backup of a particular disk to terminate immediately, but FDR may continue and attempt to backup other disks requested in the same step.
- Some errors are non-critical and the messages are warnings only. FDR will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with a U0888 abend to call your attention to the errors. Remember that a U0888 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer not to get a U0888 abend on a non-terminating error, the FDRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice ([see Section 90](#)).

CONTINUED . . .

10.02 CONTINUED**TAPE VTOC
LISTING**

Utility program FDRABRP, the ABR report program, is distributed to customers who are licensed only for FDR. Even if ABR is not licensed, you can still execute the FDRABRP "PRINT TVTOC" function. PRINT TVTOC (Tape VTOC) will list information about all of the data sets on a FDR, DSF, ABR or SAR backup (tape or disk); this can be useful when locating backups or pre-allocating data sets. This listing can be in IEHLIST LISTVTOC format or ABR VTOC format; for ICF VSAM clusters, a simulated IDCAMS LISTCAT report can also be produced. PRINT TVTOC requires only the control data (extracted VTOC and VVDS records) recorded at the very beginning of the backup, so it takes only a few seconds. Details can be found in [Sections 53.02](#) (FDRABRP JCL), 53.09 (PRINT TVTOC command), and 53.10 (Example). Other functions of FDRABRP are disabled for non-ABR customers.

FDREPORT, the FDR generalized reporting program, is also capable of reporting on the DASD data sets in a FDR backup data set. FDREPORT allows you to report on hundreds of attributes of those data sets in a format you specify. FDREPORT is not included with FDR-only licenses, but if you are licensed for ABR, or you have licensed FDREPORT as a separate product, it can be used on backups created by all FDR components.

10.03 FDR JCL REQUIREMENTS

To execute FDR, the following JCL statements are required.

**STEPLIB or
JOBLIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDR resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDR), region requirement (REGION=, [see Section 10.02](#)), and optional PARM= operand. The PARM options are a shortcut way of specifying the FDR operation to be performed; PARM values are:

1. no PARM -- FDR will process as specified by the DUMP, RESTORE or COPY command in the SYSIN data set. If there is no PARM field and no SYSIN data set, FDR will act as if PARM=D was specified; however, there is an option in the FDR Global Option Table to cause a failure if no PARM and no SYSIN are present, to avoid unintentional backups.
2. PARM=D -- FDR will dump all disks for which a pair of DISKx, TAPEx DD statements appear. Dumping will occur sequentially, in DISK DD statement order. This is the Single Disk Mode or Serial Mode described in [Section 10.02](#).
3. PARM=A -- FDR will dump concurrently all disks for which a triplet of DISKx, TAPEx, SYSPRINx DD statements appear. This is the same as ATTACH option on the DUMP statement ([Section 10.04](#)). FDR will serialize dumps to TAPEx DD cards that specify the same tape drive. A maximum of thirty TAPEx DDs may be specified but they may not point to more than 9 unique tape units. This is the Attach Mode described in [Section 10.02](#).
4. PARM=R -- FDR will restore all DISKx and TAPEx pairs, copying the volume serial from the backup.
5. PARM=N -- FDR will restore all DISKx and TAPEx pairs, retaining the volume serial of the output volume.

The PARM field may also contain an FDR DUMP, RESTORE, or COPY statement, for example,

```
PARM=' DUMP TYPE=FDR,DSNENQ=USE '
```

It will be processed as if read from SYSIN; no SYSIN DD is required.

If FDR is invoked from another program, you can pass control statements using IBM's convention for passing data from the PARM field; contact Innovation for assistance if you wish to do this.

**SYSPRINT DD
STATEMENT**

Specifies the output message data set. It must be present and is usually a SYSOUT data set but it may be assigned to disk or tape. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**SYSPRINx DD
STATEMENT**

Specifies the output data set for messages related to the matching DISKx when the ATTACH option, MAXTASKS= option or COPY statement is used (see SYSPRINT for details). Must be present if PARM=A, ATTACH, MAXTASKS= or COPY is used but it is not used without them. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD.

**FDRSUMM DD
STATEMENT**

(Optional) if present, FDR will write one-line messages for each volume dumped or restored, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set. For backups, FDRSUMM is used only if RTC=YES or DCT=YES is specified on the DUMP statement.

**SYSUDUMP DD
STATEMENT**

Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

10.03 CONTINUED

**DISKx DD
STATEMENT**

For DUMP or COPY, specifies the input disk volume. For RESTORE, specifies the disk volume whose contents will be replaced.

The format will be:

```
//DISK1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial assigned of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). Only a single disk volume serial may be specified. You may use either DISP=OLD or DISP=SHR; it makes no difference.

"x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character and must have a corresponding TAPEx statement, so there can be a maximum of 39 DISKx DDs. Processing will proceed for as many pairs of DISKx/TAPEx statements as are present, in the order that the DISKx DDs are present in the FDR JCL. If DUMMY is specified, this DD statement will be ignored.

For FDR InstantBackup

If you are also licensed for FDR InstantBackup, you can direct FDR to backup an offline point-in-time image of the volume to be backed up. This allows you to capture that point-in-time image and back it up at your leisure, even while updates are being done to the live, online disk. This may require special options on a DISKx DD statement. FDR InstantBackup is described in [Sections 25 through 29](#), with separate sections for each hardware platform on which FDR InstantBackup is supported.

10.03 CONTINUED

**TAPEX DD
STATEMENT**

Used to specify the output data set for DUMP, the input data set for RESTORE and the output disk volume for COPY. "x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character. Multiple TAPEX DD statements may be present in the FDR step JCL; a unique value for 'x' must be used for each of them (e.g., TAPE1, TAPE2, etc.). There must be a TAPEX for each DISKx; if you code one or the other, not both, it will be ignored.

For DUMP Operations:

Specifies a tape or sequential disk data set to which the backup will be written. You **must** provide a TAPEX DD statement for each volume to be backed up in this step; the TAPEX DD will receive the backup of the volume specified by DISKx. If PARM=A or ATTACH is specified, FDR will attempt to attach that many concurrent backup subtasks, but may postpone some of them if it detects that they require a tape drive in use by another backup (see the notes on UNIT=AFF and VOL=REF below). Up to 39 TAPEX DDs may be present but they must not point to more than 9 unique tape units.

DUMMY is supported, for **testing** purposes only. The disk will be read, but the backup data will be discarded.

You must provide all the JCL parameters required to allocate and catalog the backup data set on disk or tape, which may include some or all of: DSN=, UNIT=, VOL=, SPACE=, and DISP=(NEW,CATLG). For tape, a volume count should be specified since the default is only 5 tape volumes, e.g., VOL=(,,,255).

DCB parameters are not required and should be omitted. However, tape unit hardware compaction (sometimes called IDRC), available on most tape cartridge drives) can be requested by adding DCB=TRTCH=COMP to your DD statement; Tape hardware compaction may be the default depending on local MVS options. For tapes attached by ESCON or FICON channels, Innovation recommends use of tape hardware compaction instead of FDR software compression (the COMPRESS= option).

```
Examples: //TAPE1 DD DSN=PROD.MVS001.BACKUP1,UNIT=TAPE,
           //          VOL=(, , , 255),DISP=(NEW,CATLG)
           //TAPE2 DD DSN=PROD.MVS002.BACKUP2,UNIT=DISK,
           //          SPACE=(CYL,(500,100),RLSE),DISP=(,CATLG)
```

For tape backups, UNIT=AFF or VOL=REF may be specified, referencing another TAPEX DD statement, to reduce the number of tape drives used in the step. UNIT=AFF=TAPEX will cause MVS to allocate the same tape drive for both DD statements, but will call for separate output tapes when each DD is opened. VOL=REF=*.TAPEX with LABEL=n can be used to stack multiple backup files on the same tape, providing more complete utilization of the tape volumes (which may be important as new technology increases tape volume capacity). FDR will automatically recognize that multiple TAPEX DDs point to the same tape drive and will serialize operations on that drive so that only one backup is directed to that drive at a time.

```
Examples: /* The following creates 2 backups on 2 different
           /* tape volumes using the same tape drive. This may not
           /* fully utilize the tape volumes but will allow for
           /* concurrent restores from these backups.
           //TAPE1 DD DSN=PROD.MVS001.BACKUP1,UNIT=TAPE,
           //          VOL=(, , , 255),DISP=(NEW,CATLG)
           //TAPE2 DD DSN=PROD.MVS002.BACKUP2,UNIT=AFF=TAPE1,
           //          VOL=(, , , 255),DISP=(NEW,CATLG)

           /* The following creates a multi-file (and possibly
           /* multi-volume) tape containing 3 backups.
           //TAPE3 DD DSN=PROD.MVS003.BACKUP3,UNIT=TAPE,
           //          VOL=(,RETAIN,,255),DISP=(NEW,CATLG)
           //TAPE4 DD DSN=PROD.MVS004.BACKUP4,LABEL=2,
           //          VOL=(,RETAIN,REF=*.TAPE3),DISP=(NEW,CATLG)
           //TAPE5 DD DSN=PROD.MVS005.BACKUP5,LABEL=3,
           //          VOL=(,RETAIN,REF=*.TAPE4),DISP=(NEW,CATLG)
```

CONTINUED ...

10.03 CONTINUED

For DUMP Operations: (continued)

Your tape management software may require that you add an operand to the TAPEX DD to specify when the tape will be returned to the scratch pool. The operands are:

RETPD=nnnn or **LABEL=RETPD=nnnn** retain tape for "nnnn" days.

EXPDT=yyyyddd or **LABEL= EXPDT=yyyyddd** retain tape until Julian date yyyy.ddd

Certain expiration date values are treated as keywords by some tape management systems, e.g., EXPDT=99000 may indicate "expire when the backup data set is no longer cataloged". See [Section 80.32](#) for details.

If you are running the StorageTek ExHPDM (High Performance Data Mover) software product, you can direct FDR backups to ExHPDM with the SUBSYS= JCL operand, e.g.,

```
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,DISP=(NEW,CATLG),
// SUBSYS=(SOV,'CLASS(FDRBKUP)')
```

Please read [Section 80.33](#) and the ExHPDM program documentation for more details.

Backups can also be written to a sequential data set on disk, e.g.,

```
//TAPE1 DD DSN=TECH.BACKUP1,UNIT=3390,VOL=SER=TECH01,
// SPACE=(CYL,(25,5),RLSE),DISP=(,CATLG)
```

Although the backup will usually take less space than the original data being backed up, this can vary, so be sure to specify a secondary allocation quantity and the RLSE parameter to release unused cylinders.

For RESTORE Operations:

Specifies a backup data set on tape or disk from which the data is to be restored. The backup must, of course, be a full-volume backup created by FDR, ABR or SAR. You may include multiple TAPEX DD statements in order to restore multiple volumes. The corresponding DISKx DD statement is required to define which backup is to be restored to which volume.

Example: //TAPE1 DD DSN=PROD.BACKUP1,DISP=OLD

If the backup data set is not cataloged, you must specify a valid UNIT= and VOL=SER= specifying all volumes containing the backup, in the right order. LABEL=n may also be required if the backup is not file 1 on the tape.

Example: //TAPE1 DD DSN=PROD.BACKUP1,DISP=OLD,
// UNIT=CART,VOL=SER=B01279,LABEL=3

If you provide multiple TAPEX DD statements, MVS will allocate a separate tape unit for each DD (except for DDs that point to the same tape volser). If you are using the default of MAXTASKS=1, FDR will restore those backups serially (one at a time) so it unnecessarily ties up the tape units until all restores are done. You can minimize this waste by using UNIT=AFF to mount all the backups on a single tape unit.

Examples: /* The following mounts 3 backups on the same tape drive.

```
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,DISP=OLD
//TAPE2 DD DSN=PROD.MVS002.BACKUP1,UNIT=AFF=TAPE1,DISP=OLD
//TAPE3 DD DSN=PROD.MVS003.BACKUP1,UNIT=AFF=TAPE1,DISP=OLD
```

Even if you specify MAXTASKS=, FDR will automatically recognize that multiple TAPEX DDs point to the same tape drive and will serialize operations on that drive so that only one restore is executed on that drive at a time.

If the backup to be restored was directed to the StorageTek ExHPDM software product as described earlier, you must specify the SUBSYS= operand on the TAPE DD statement to invoke ExHPDM and restore that backup, e.g.,

```
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,DISP=OLD,SUBSYS=SOV
```

10.03 CONTINUED

For COPY Operations:

Specifies the output disk volume and unit using the same format used for DISKx

Example: //TAPE1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD

**TAPExx DD
STATEMENT**

Specifies a second backup data set, using the same format documented for TAPEx (DUMP operation) above. A backup identical to TAPEx will be produced on TAPExx; the same data blocks are written to both simultaneously. For example, if DISK6 is being dumped to TAPE6, the inclusion of a TAPE66 DD statement will cause a second backup file to be produced.

On COPY operations, this is the only sequential backup of the DISKx volume (since TAPEx defines the output disk volume).

TAPExx is optional; if omitted, no duplicate backup is created; for COPY, no backup is created at all.

**SYSIN DD
STATEMENT**

Optional control statement data set. Usually an input stream or DD * data set. If SYSIN is allocated with a disposition of NEW, FDR will ignore this data set, so the data set must either be an input stream data set or an existing disk data set. If the control statement was provided in the EXEC PARM=, it can be DUMMY.

Warning: if you omit SYSIN or omit the DISP= parameter for a disk data set (which assumes NEW), FDR will use only the PARM= option to define the operation to be performed. If that is also omitted, it assumes D (DUMP); this may result in a backup being taken when you intended a restore, overlaying the backups you wanted to restore. There is an option in the FDR Global Option Table (on ISPF panel A.I.4.2) which causes FDR to fail if SYSIN is ignored.

CONTINUED ...

10.04 FDR DUMP STATEMENT

DUMP	TYPE=FDR	,ENQ=ON OFF RESERVE
D	,ATTACH MAXTASKS=n	,ENQERR=NO
	,BUFNO=MAX nn	,FORMAT=NEW SPLIT
	,COMPRESS=ALL COPY1 COPY2	,HFS=QUIESCE
	,DATA=ALL USED	,MAXERR=nnnn
	,DCT=YES NO	,RTC=YES NO
	,DSNENQ=NONE TEST USE HAVE	,SNAP=(USE,REL)

**DUMP
STATEMENT**

The DUMP statement requests a full-volume dump (backup) operation. It is optional (see PARM= in [Section 10.03](#)) but if present must be the first and only statement input. Only one DUMP, RESTORE or COPY statement is allowed per execution.

These operands will control the backup of each specified disk (DISKx DD statement) to its backup data set (TAPEEx DD statement). If a TAPEEx DD is also present, a second duplicate copy of the backup is written there. The backup data sets may be on tape or disk. The backups will proceed for each valid DISKx/TAPEEx pair in the JCL. If the ATTACH or MAXTASKS= option is specified, a SYSPRINx DD is also required, to separate messages for each concurrent dump subtask.

OPERANDS	TYPE=FDR	Specifies that a full volume dump is to be performed. It must be specified on the DUMP Statement.
	ATTACH	Backups will proceed concurrently for all disks for which a triplet of DISKx, TAPEEx, SYSPRINx DD statements appear. However, backups to TAPEEx DD cards that specify the same tape drive will be automatically serialized. Up to 39 TAPEEx DDs may be specified, but they must not point to more than 9 unique tape units (use UNIT=AFF= or VOL=REF= to reduce the number of allocated drives). This is the same function as specified by using PARM=A as shown in Section 10.03 . Default: each DISKx/TAPEEx pair will be processed one at a time, in the order that the DISKx DDs appear in the JCL, unless MAXTASKS= is specified.
	BUFNO=	specifies how many buffers will be used for dumping each disk volume. Each buffer holds one disk track. The buffers acquired will be divided into 2 sets in order to overlap input and output I/O operations; each disk I/O will read disk tracks into one half of the buffers. MAX – buffers sufficient to read 1 cylinder of the input disk are acquired. nn - the specified number of buffers is acquired. Default: MAX unless overridden in the FDR Global Option Table. Innovation recommends that you do not override the default. However, BUFNO=2 will be forced when a backup (output) data set is on disk.

10.04 CONTINUED

- COMPRESS=** Controls the use of FDR software compression. Values for COMPRESS= are:
ALL -- the backup file for both copies (TAPEX and TAPEXX) is to be compressed.
COPY1 -- only the backups on TAPEX DD statements will be compressed.
COPY2 -- only the backups on TAPEXX DD statements will be compressed.
See "Memory Requirements" in [Section 10.02](#) for the additional storage required by COMPRESS=. COMPRESS= is ignored if DCT=YES is also specified.

Default: backups will not be compressed.

COMPRESS is recommended for backups to disk files, and for tape backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of tape hardware compaction, e.g., IDRC, usually results in better performance.
- DATA=** **USED** -- only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be backed up. On most volumes, this will make the dump run faster.
ALL -- all allocated tracks of all data sets will be backed up. You may need to specify DATA=ALL if the data sets to be backed up include JES2 spool data sets or CICS log data sets, since they usually do not have valid last block pointers.

Default: ALL.
- DCT=** DCT= is valid only if you are licensed for FDR InstantBackup. It will be honored only if the disk being backed up is in a disk subsystem with the HSDM option (High Speed Data Mover). HSDM allows FDR to backup and restore the internal compressed images of disk tracks, improving backup elapsed times up to 60%. It can also be specified as DUMPCOMPRESSEDTRACK=.
- YES** -- use HSDM for any volume where the disk hardware has the HSDM feature installed. Normal backup will be used for other volumes.
NO -- do not use HSDM.

Default is NO.

Note that DCT=YES implies RTC=YES; see the description of RTC= for its benefits.

CONTINUED . . .

10.04 CONTINUED

DSNENQ=

Specifies whether all of the data sets on the volume being dumped will be ENQed. See "Data Set Enqueue Option" in [Section 10.02](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a FDR158 warning message is issued for the data set but the data set will still be dumped since this is a full-volume backup. A successful ENQ will prevent any other task from using the data set until the backup of that volume is complete. The options for DSNENQ= are:

USE -- The data sets will be enqueued for the duration of the backup from this disk volume. This is the most frequently used option.

TEST -- The data sets will only be tested to see if they are enqueued to another task at the time that the dump from this volume starts.

HAVE -- The data sets will be enqueued for the duration of the backup from this disk volume. If not available, a message (FDRW27) is issued to the MVS operator. See "Data Set Enqueue Option" in [section 10.02](#) for the valid responses.

NONE -- No data set ENQ will be issued.

CAUTION: *This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.*

Default: NONE.

Recommendation: use DSNENQ=USE if you want to be sure that no other task uses the input data sets until the backup is complete.

ENQ=

Specifies whether an ENQ should be done on the VTOC of each disk volume while it is being backed up. See "VTOC Enqueue Option" in [Section 10.02](#) for more details.

ON -- the VTOC of each disk volume will be ENQed during its backup. This ENQ may be effective only on the system where the backup is executing; other systems may still be able to update the VTOC.

RESERVE -- in addition to the ENQ, a hardware RESERVE will be issued on each disk volume during its backup. This is meaningful only on a system with "shared DASD" where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF -- the VTOC will not be enqueued or reserved during the backup.

Default: OFF.

10.04 CONTINUED

- ENQERR=** **NO** – If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see “Step Termination” in [Section 10.02](#)). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally.
- Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.
- FORMAT=** Specifies the format of the sequential backup file.
- NEW** -- the backup will be created using a maximum of a 56K blocksize. A block will contain the image of one or more tracks from the input disk.
- SPLIT** -- the backup will be created using a maximum blocksize of 32K. For blocks that would be more than 32K, they are written as 2 blocks of 32K or less. However, FORMAT=SPLIT causes use of a disk I/O technique which is less efficient than that used by FORMAT=NEW, which will impact backup performance.
- WARNING: If you use a normal copy program (ex: IEBGENER) to copy a backup file created with FORMAT=NEW, you will not get any error messages, but the resulting tape will not be usable for a restore. Tapes in the new format must only be copied with the INNOVATION provided tape copy program (FDRTCOPY) or FATAR.***
- Default: NEW if backup on tape -- SPLIT if backup is on disk.
- HFS=** **QUIESCE** invokes special processing when HFS (Hierarchical File System) data sets are backed up. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to Unix System Services, so a "quiesce" call is issued to prevent updates to the HFS data set during the backup. See "HFS" in [Section 80.11](#) for details and security requirements.
- Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for **all** data sets being backed up, not just HFS data sets.
- Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table ([see Section 90](#)). If you use the default, you should unmount the HFS file system before the backup to be sure of getting a usable backup.
- MAXERR=** Specifies the number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup is written to the ExHPDM subsystem, MAXERR=1 is recommended.
- WARNING: MAXERR over the default value may result in the loss of many data sets. This option should only be used when necessary and with care.***
- Default: 20 errors.

10.04 CONTINUED

MAXTASKS=

Specifies the maximum number of volumes that will be dumped concurrently in this step. It has the same affect as the ATTACH operand except that the maximum number of concurrent backups is limited to "n". The value may be from 1 to 9 but the actual number of concurrent backups will be no greater than the number of TAPEX DD statements in this step's JCL (if multiple TAPEX DD statements point to the same tape drive, one only concurrent dump will use that drive at any time). If MAXTASKS=2 or more, messages for each backup are directed to the SYSPRINx DD corresponding to the TAPEX DD associated with the backup ([see Section 10.03](#)).

The default is 1 unless the ATTACH operand is specified.

RCT=

It can also be specified as READTRACKCCW=.

YES -- use READ TRACK CCWs to read disk data tracks. RTC=YES also causes:

- up to 1 cylinder of disk data is read at a time.
- FDR buffers are moved above the 16MB line (about 2MB per concurrent backup), allowing more concurrent backups to be run in one step.
- the elapsed time of FDR backups when the backup data set is itself on disk is significantly improved.

NO -- use other techniques to read disk data tracks.

Default is NO.

SNAP=

Used only when you are licensed for FDR InstantBackup and only if you have previously used program FDRSNAP to create an instant point-in-time image of a volume in an IBM RVA or StorageTek Iceberg/SVA disk subsystem with the Snapshot feature. [See Section 26 for details of the use of SNAP= with FDR.](#)

USE - tells FDR InstantBackup to read the offline snapped copy of the online volume specified in JCL. FDR remembers the device address of the offline device most recently used as a target for a snap of each online volume unless an intervening IPL has occurred.

(USE,REL) - same as SNAP=USE, but at the end of backing up each snapped volume, FDR will issue a request to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the disk subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.

Default: backup the online disk for each volume.

10.05 FDR RESTORE STATEMENT

```

RESTORE          TYPE=FDR              ,MAXERR=nnnnn
R               ,CONFMESS=YES|NO      ,MAXTASKS=n
               ,CPYVOLID=NO|YES       ,PROT=NONE
               ,ENQ=OFF|RESERVE      ,SMSPROT=ALL|NONE
                                   ,VOLRESET=NO|YES

```

**RESTORE
STATEMENT**

The RESTORE statement requests a full-volume restore operation. It is optional (see PARM= in [Section 10.03](#)) but if present must be the first and only statement input. Only one DUMP, RESTORE or COPY statement is allowed per execution.

These operands will control the restore of each specified backup data set (TAPEx DD statement) to its output disk (DISKx DD statement). The backup must be a full-volume backup created by FDR, ABR, or SAR.

A FDR RESTORE operation can restore from a backup of a DASD volume to a larger capacity volume of the same DASD type (e.g., 3380-E to 3380-K or 3390-2 to 3390-3). If you want to restore from a larger DASD (such as 3390-3) to a smaller device (3390-2) it is usually more convenient to use COMPAKTOR ([Section 40](#)) or FDRDSF ([Section 20](#)). However, if you prefer to use RESTORE TYPE=FDR and no data sets are allocated on tracks beyond the end of the smaller disk, you can do so by specifying PROT=NONE.

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

OPERANDS	TYPE=FDR	Specifies that a full volume restore is to be performed. It must be specified on the RESTORE statement.
	CONFMESS=	<p>YES -- before beginning the restore, FDR will request confirmation via a WTOR (FDRW01) message to which the MVS operator must reply.</p> <p>NO -- suppresses the WTOR and begins the restore immediately.</p> <p>Default: YES.</p> <p>NOTE: CONFMESS=NO can be very useful at a disaster recovery site to avoid full volume restores being delayed waiting for an operator response.</p>

10.05 CONTINUED

CPYVOLID=

Specifies whether the volume serial number of the disk that was backed up will be restored, if the existing volser of the output disk is different (if the serials are the same, CPYVOLID is ignored).

YES -- volume serial number of the output volume will be replaced with the original volume serial number of the disk which was dumped. This option can also be specified by PARM=R on the JCL EXEC statement. If another online volume has the same serial, the restored volume will be placed offline at the end of the restore.

NO -- the volume serial number of the output volume will be retained. This option can also be specified by PARM=N on the JCL EXEC statement. See the VOLRESET= operand below.

Default: NO -- unless the volume being restored was SMS-managed, when YES is forced.

Note: although full-volume FDR restore does not catalog any data sets, any data sets which were cataloged to the original volume are automatically cataloged to the new volume when restoring with CPYVOLID=YES. If you use CPYVOLID=NO and do not later relabel the volume, data sets may need to be manually recataloged.

VSAM/SMS WARNING: See the notes under VOLRESET=NO below.
CPYVOLID=YES is recommended for any volume containing a VVDS.

ENQ=

Specifies whether an ENQ should be done on the VTOC of each disk volume during the restore. See "VTOC Enqueue Option" in [Section 10.02](#) for more details.

RESERVE --a hardware RESERVE will be issued on each disk volume during its restore. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF -- the VTOC will not be enqueued or reserved during the restore.

Default: RESERVE.

MAXERR=

Specifies the number of disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup was written to the ExHPDM subsystem, MAXERR=1 is recommended.

WARNING: MAXERR over the default value may result in the loss of many data sets. This option should only be used when necessary and with care.

Default: 20 errors.

10.05 CONTINUED

MAXTASKS= Specifies the maximum number of volumes that will be restored concurrently in this step. The value may be from 1 to 9 but the actual number of concurrent restores will be no greater than the number of TAPEX DD statements in this step's JCL (if multiple TAPEX DD statements point to the same tape drive, one only concurrent restore will use that drive at any time).

The default is 1. However, if MAXTASKS= is omitted all restore messages will be printed on SYSPRINT while if MAXTASKS=n is specified (even if "n" is 1), messages are directed to the SYSPRINx DD corresponding to the TAPEX DD associated with the restore ([see Section 10.03](#)).

PROT= **NONE** -- FDR can only restore to volumes of the same type as the dumped disk and normally requires that the output disk have the same or larger capacity (number of cylinders) as the input disk, e.g., 3390-2 to 3390-2 or 3390-3. PROT=NONE allows FDR to restore to a smaller disk of the same type, e.g., 3390-3 to 3390-2. Do not specify PROT=NONE unless you need this function since it also suppresses other validity checks.

WARNING: Since FDR always places data tracks in their original locations, restore to a smaller disk will fail if the larger disk had data sets allocated beyond the bounds of the smaller disk. Use COMPAKTOR to convert such disks.

Default: all validity checks are done.

SMSPROT= **ALL** -- enforces several rules when SMS-managed volumes are involved: Backups of SMS-managed volumes can only be restored to SMS-managed volumes, and non-SMS volumes only to non-SMS volumes. CPYVOLID=YES is forced when an SMS-managed volume is restored.

NONE -- allows the restore of SMS-managed volumes to non-SMS volumes, and vice versa. Also allows the restore of SMS volumes to new volumes if CPYVOLID=NO is specified.

WARNING: SMSPROT=NONE should be used with caution. It will usually be used at a disaster recovery site where a relPL of MVS will be done after the restores, to place all volumes in the proper SMS status. See Section 70 for more details on restoring and moving SMS-managed volumes.

Default: ALL.

CONTINUED . . .

10.05 CONTINUED

VOLRESET=

When CPYVOLID=NO or PARM=N is specified or defaulted, and the volume serial of the output disk is different from that of the original disk on the backup data set, the backup is restored but the volume serial of the output volume is retained. VOLRESET= controls additional processing relating to this change in volume serials. VOLRESET is ignored if the volume serial of the output disk is not being changed.

YES -- the volume serials that are part of the data set names of the VTOC Index ("SYS1.VTOCIX.volser") and the ABR model DSCB (usually "FDRABR.Vvolser") are checked to see if they match the input volume serial (the volume backed up). If so, they are renamed to match the volser of the output disk. Also, the DSCB field DS1DSSN (data set serial number, usually the volser of the first or only volume of the data set) for every data set on the volume will be changed to the new volume serial if the existing value matched the original volume serial. VOLRESET=YES should be used with CPYVOLID=NO when you intend to permanently retain the volume serial of the output disk. Note that data sets on the volume will not be recataloged to the new volume serial.

NO -- do not rename the VTOCIX and ABR model DSCB, and do not change DS1DSSN fields. VOLRESET=NO should be used with CPYVOLID=NO if you plan to eventually relabel the disk back to the original volume serial.

NOTE: If a volume initialized for ABR is restored, FDR will reset the ABR model DSCB to force a full volume dump on the next execution of ABR volume backup, unless VOLRESET=NO is specified.

VSAM/SMS WARNING: FDR will not rename the VVDS, since the VVDS and the catalogs contain self-defining records that would also need resetting. If a volume containing VSAM clusters or SMS-managed data sets is restored with CPYVOLID=NO, the data sets will be inaccessible unless the volume is relabeled to the original serial number.

CPYVOLID=YES is recommended for any volume containing a VVDS, but if it is necessary to restore the volume under a temporary volser, use CPYVOLID=NO and VOLRESET=NO and relabel the volume with ICKDSF later.

Default: YES but it is ignored if CPYVOLID=YES (or PARM=R) is also specified.

10.06 FDR COPY STATEMENT

COPY	TYPE=FDR	,ENQ=ON OFF RESERVE
	,COMPRESS=COPY2	,HFS=QUIESCE
	,CONFMESS=YES NO	,MAXERR=nnnnnn
	,CPYVOLID=NO YES	,PROT=NONE
	,DATA=ALL USED	,SMSPROT=ALL NONE
	,DSNENQ=NONE TEST USE HAVE	,VOLRESET=NO YES

**COPY
STATEMENT**

The COPY statement requests a full-volume disk-to-disk copy. It is required if a copy is to be performed and must be the first and only statement input. Only one DUMP, RESTORE or COPY statement is allowed per execution. Each disk volume specified by a DISKx DD statement will be copied to the disk volume specified by the matching TAPEx DD statement; TAPEx must point to the same type of disk (e.g., 3390) as DISKx. If a TAPExx DD statement is also specified, FDR will also create a backup of DISKx.

A COPY is essentially a DUMP and a RESTORE in separate subtasks under FDR, so most options documented for DUMP and RESTORE in the preceding sections apply. The backup data created by the DUMP subtask is passed directly to the RESTORE subtask. No backup data set is created unless you provide a TAPExx DD statement to create a backup as well as doing the copy (see [Section 10.03](#)).

NOTE: A SYSPRINx DD statement is required matching DISKx/TAPEx pair. This print data set will record the messages from the dump of the DISKx volume. The SYSPRINT data set will record the messages from the restore to the TAPEx volume.

A FDR COPY operation can copy from a DASD volume to a larger capacity volume of the same DASD type (e.g., 3380-E to 3380-K or 3390-2 to 3390-3). If you want to copy from a larger DASD (such as 3390-3) to a smaller device (3390-2) it is usually more convenient to use COMPAKTOR ([Section 40](#)) or FDRCOPY ([Section 21](#)). However, if you prefer to use COPY TYPE=FDR and no data sets are allocated on tracks beyond the end of the smaller disk, you can do so by specifying PROT=NONE.

WARNING: before you do a full-volume copy you should make sure that the target volume is offline to all systems other than the system where the copy is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

OPERANDS **TYPE=FDR** Indicates that the entire volume is to be copied. Must be specified.

10.06 CONTINUED

COMPRESS= Controls the use of FDR software compression. COMPRESS is valid only if you have a TAPExx DD statement to create a backup of the input disk while copying. Values for COMPRESS= are:

COPY2 -- the backups on TAPExx DD statements will be compressed.
Add 1024K (1M) to the memory requirement when COMPRESS= is specified.

Default: backups will not be compressed..

COMPRESS is recommended for backups to disk files, and for tape backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of tape hardware compaction, e.g., IDRC, usually results in better performance.

CONFMESS= **YES** -- before beginning the copy, FDR will request confirmation via a WTOR message (FDRW01) to which the MVS operator must reply.
NO -- suppress the WTOR and begin the copy immediately.

Default: YES.

CPYVOLID= Specifies whether the volume serial number of the input disk will be copied to the output disk
YES -- volume serial number of the output volume will be replaced with the volume serial number of the input disk. The output volume will be placed offline at the end of the copy.

NO -- the volume serial number of the output volume will be retained. See the VOLRESET= operand below.

Default: NO -- unless the volume being copied was SMS-managed, when YES is forced.

Note: although full-volume FDR copy does not catalog any data sets, any data sets which were cataloged to the original volume are automatically cataloged to the new volume when copying with CPYVOLID=YES. If you use CPYVOLID=NO and do not later relabel the volume, data sets may need to be manually recataloged.

VSAM/SMS WARNING: See the notes under VOLRESET=NO below.
CPYVOLID=YES is recommended for any volume containing a VVDS.

DATA= **USED** -- only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be copied. On most volumes, this will make the copy run faster.

ALL -- all allocated tracks of all data sets will be copied.

Default: ALL.

10.06 CONTINUED

DSNENQ=

Specifies whether all of the data sets on the input volume will be ENQed. See “Data Set Enqueue Option” in [Section 10.02](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a warning message is issued for the data set but the data set will still be copied since this is a full-volume copy. A successful ENQ will prevent any other task from using the data set until the copy of that volume is complete. The options for DSNENQ= are:

USE -- The data sets will be enqueued for the duration of the copy from this disk volume. This is the most frequently used option.

TEST -- The data sets will only be tested to see if they are enqueued to another task at the time that the copy from this volume starts.

HAVE -- The data sets will be enqueued for the duration of the copy. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not enqueue the data set)

RETRY (try the enqueue again)

NONE -- No data set ENQ will be issued.

CAUTION: *This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.*

Default: NONE.

ENQ=

Specifies whether an ENQ should be done on the VTOC of each input and output disk volume during the copy. See “VTOC Enqueue Option” in [Section 10.02](#) for more details.

ON -- the VTOC of only the disk volume will be ENQed during the copy. This ENQ may be effective only on the system where the backup is executing; other systems may still be able to update the VTOC.

RESERVE -- an ENQ and a hardware RESERVE will be issued on both the input and output volumes during the copy. This is meaningful only on a system with “shared DASD” where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF -- the VTOC will not be enqueued or reserved during the backup.

Default: OFF (no ENQ) on the input volume and RESERVE on the output volume.

CONTINUED . . .

10.06 CONTINUED

HFS=

QUIESCE invokes special processing when HFS (Hierarchical File System) data sets are copied. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to Unix System Services, so a "quiesce" call is issued to prevent updates to the HFS data set during the backup. See "HFS" in [Section 80.11](#) for details and security requirements.

Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for **all** data sets being copied, not just HFS data sets.

Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table ([see Section 90](#)). If you use the default, you should unmount the HFS file system before the backup to be sure of getting a usable copy.

MAXERR=

Specifies the number of disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP.

WARNING: MAXERR over the default value may result in the loss of many data sets. This option should only be used when necessary and with care.

Default: 20 errors.

PROT=

NONE -- FDR can only copy to volumes of the same type as the dumped disk and normally requires that the output disk have the same or larger capacity (number of cylinders) as the input disk, e.g., 3390-2 to 3390-2 or 3390-3. PROT=NONE allows FDR to copy to a smaller disk of the same type, e.g., 3390-3 to 3390-2. Do not specify PROT=NONE unless you need this function since it also suppresses other validity checks.

WARNING: Since FDR always places data tracks in their original locations, copy to a smaller disk will fail if the larger disk had data sets allocated beyond the bounds of the smaller disk. Use COMPAKTOR to convert such disks.

Default: all validity checks are done.

SMSPROT=

ALL -- enforces several rules when SMS-managed volumes are involved: SMS-managed volumes can only be copied to SMS-managed volumes, and non-SMS volumes only to non-SMS volumes. CPYVOLID=YES is forced when an SMS-managed volume is copied.

NONE -- allows the copy of SMS-managed volumes to non-SMS volumes, and vice versa. Also allows the copy of SMS volumes to new volumes if CPYVOLID=NO is specified.

WARNING: SMSPROT=NONE should be used with caution. If used, it may be necessary to reIPL or restart SMS before the new volume location is recognized. See Section 70 for more details on restoring and moving SMS-managed volumes.

Default: ALL.

10.06 CONTINUED

VOLRESET=

When CPYVOLID=NO is specified or defaulted, the input disk is copied but the volume serial of the output disk is retained. VOLRESET= controls additional processing relating to this change in volume serials.

YES -- the volume serials that are part of the data set names of the VTOC Index ("SYS1.VTOCIX.volser") and the ABR model DSCB (usually "FDRABR.Vvolser") are checked to see if they match the input volume. If so, they are renamed to match the volser of the output disk. Also, the DSCB field DS1DSSN (data set serial number, usually the volser of the first or only volume of the data set) for every data set on the volume will be changed to the new volume serial if the existing value matched the original volume serial.

VOLRESET=YES should be used with CPYVOLID=NO when you intend to permanently retain the volume serial of the output disk. Note that data sets on the volume will not be recataloged to the new volume serial.

NO -- do not rename the VTOCIX and ABR model DSCB, and do not change DS1DSSN fields. VOLRESET=NO should be used with CPYVOLID=NO if you plan to eventually place the original disk offline and relabel the output disk back to the input volume serial.

NOTE: If a volume initialized for ABR is copied, FDR will reset the ABR model DSCB to force a full volume dump on the next execution of ABR volume backup, unless VOLRESET=NO is specified.

VSAM/SMS WARNING: *FDR will not rename the VVDS, since the VVDS and the catalogs contain self-defining records that would also need resetting. If a volume containing VSAM clusters or SMS-managed data sets is copied with CPYVOLID=NO, the data sets will be inaccessible unless the volume is relabeled to the original serial number.*

CPYVOLID=YES is recommended for any volume containing a VVDS, but if it is necessary to copy the volume under a temporary volser, use CPYVOLID=NO and VOLRESET=NO and relabel the volume with ICKDSF later.

Default: NO. By default, COPY does not rename these data sets.

CONTINUED . . .

10.10 FDR DUMP EXAMPLES

The examples in the [sections 10.10-10.12](#) illustrate usage of FDR and some of the FDR options. JOB statements are not shown. If FDR is not in the system linklist, you will need to add a JOBLIB or STEPLIB DD statement pointing to the FDR program library.

Note that FDR determines the type of disk being processed from the UCB for the disk device, not from the UNIT= JCL parameter. Any UNIT= value which will allocate the disk volumes to be processed can be used.

**DUMP ONE
VOLUME**

Dump a 3380 disk volume to a GDG on 3480 tape cartridges; tape management will expire the backup when the generation is no longer cataloged. Since the backup will probably take more than one cartridge, 2 3480 tape drives are assigned via "UNIT=(3480,2)" so that the backup will not have to wait for a cartridge to be rewound and a new cartridge mounted. A volume count of 20 is provided since the backup may take more than the IBM default of 5 volumes. FDR software compression is used to improve performance of the relatively slow parallel tape channel.

```
//DUMP      EXEC   PGM=FDR,REGION=2M
//SYSPRINT  DD     SYSOUT=*
//SYSUDUMP  DD     SYSOUT=*
//DISK1     DD     UNIT=3380,DISP=OLD,VOL=SER=123456
//TAPE1     DD     UNIT=(3480,2),DSN=BACKUP.V123456(+1),
//              DISP=(,CATLG),VOL=(,,20),LABEL=EXPDT=99000
//SYSIN     DD     *
              DUMP      TYPE=FDR,COMPRESS=ALL
```

**CREATE
DUPLICATE
BACKUPS**

Dump a 3390 disk volume to tape, making two backup copies. FDR will attempt to ENQ all data sets being dumped; if the ENQ fails, a warning is issued but the data set is still backed up. The step will end with a U0888 abend if any ENQ failures occur, to draw attention to the messages; if you don't want the abend if the only error is the active datasets, add the operand "ENQERR=NO".

```
//DUMP      EXEC   PGM=FDR,REGION=2M
//SYSPRINT  DD     SYSOUT=*
//SYSUDUMP  DD     SYSOUT=*
//DISKA     DD     UNIT=3390,DISP=OLD,VOL=SER=D33901
//TAPEA     DD     UNIT=TAPE,DSN=BACKUP.COPY1,DISP=(,KEEP)
//TAPEAA    DD     UNIT=TAPE,DSN=BACKUP.COPY2,DISP=(,KEEP)
//SYSIN     DD     *
              DUMP      TYPE=FDR,DSNENQ=USE
```

**DUMP
MULTIPLE
VOLUMES
SERIALLY**

Dump three disk volumes sequentially to tape. Only one physical tape drive is used, but a fresh scratch tape will be mounted for each backup. For PS and PO data sets, only the used tracks will be dumped; for all other data set types, all allocated tracks will be dumped.

```
//DUMP      EXEC   PGM=FDR,REGION=1M
//SYSPRINT  DD     SYSOUT=*
//SYSUDUMP  DD     SYSOUT=*
//DISK3     DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0001
//TAPE3     DD     UNIT=TAPE,DSN=BACKUP.VTS0001,DISP=(,KEEP)
//DISK7     DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0002
//TAPE7     DD     DSN=BACKUP.VTS0002,DISP=(,KEEP),
//              UNIT=AFF=TAPE3
//DISK9     DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0003
//TAPE9     DD     DSN=BACKUP.VTS0003,DISP=(,KEEP),
//              UNIT=AFF=TAPE3
//SYSIN     DD     *
              DUMP      TYPE=FDR,DATA=USED
```

10.10 CONTINUED

**DUMP
MULTIPLE
VOLUMES
SERIALLY TO
ONE TAPE I**

Dump three disk volumes sequentially to tape. Only one physical tape drive is used and the backups will be piggy-backed as three consecutive files on one or more tape volumes. If your output is a high-capacity tape cartridge, such as the IBM Magstar 3590 or StorageTek Redwood or 9840, this allows you to use each tape volume more efficiently, with less wasted tape. You can dump up to 39 disk volumes, creating up to 39 tape files, in one step. DCB=TRTCH=COMP is specified to invoke tape hardware compression and further reduce the tape requirements.

```
//DUMP          EXEC   PGM=FDR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK3         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0001
//TAPE3         DD     UNIT=TAPE,DSN=BACKUP.VTS0001,DISP=(,KEEP),
//              VOL=(,RETAIN),LABEL=1
//DISK7         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0002
//TAPE7         DD     DSN=BACKUP.VTS0002,DISP=(,KEEP),
//              VOL=REF=*.TAPE3,LABEL=2,DCB=TRTCH=COMP
//DISK9         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0003
//TAPE9         DD     DSN=BACKUP.VTS0003,DISP=(,KEEP),
//              VOL=REF=*.TAPE7,LABEL=3,DCB=TRTCH=COMP
//SYSIN         DD     *
DUMP            TYPE=FDR,DATA=USED
```

**DUMP
MULTIPLE
VOLUMES
SERIALLY TO
ONE TAPE II**

If you require more than 39 backups to fill a high-capacity cartridge, you can do so with multiple steps

```
//DUMP1         EXEC   PGM=FDR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK0         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0001
//TAPE0         DD     UNIT=TAPE,DSN=BACKUP.VTS0001,DISP=(,CATLG),
//              VOL=(,RETAIN),LABEL=1
//DISK1         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0002
//TAPE1         DD     DSN=BACKUP.VTS0002,DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.TAPE0),LABEL=2
//              . . . 36 more TAPEX/DISKx pairs appear here
//DISK$         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0039
//TAPE$         DD     DSN=BACKUP.VTS0039,DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.TAPE#),LABEL=39
//SYSIN         DD     *
DUMP            TYPE=FDR,DATA=USED
//DUMP2         EXEC   PGM=FDR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK0         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0040
//TAPE0         DD     DSN=BACKUP.VTS0040,DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.DUMP1.TAPE$),LABEL=40
//DISK1         DD     UNIT=DISK,DISP=OLD,VOL=SER=TS0041
//TAPE1         DD     DSN=BACKUP.VTS0041,DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.TAPE0),LABEL=41
//              . . . additional TAPEX/DISKx pairs if required
//SYSIN         DD     *
DUMP            TYPE=FDR,DATA=USED
```

Recommendation: if you are licensed for FDRABR, ABR Volume Backups will automatically stack backups on one or more output tapes and does not require the complex JCL shown above. This will automatically maximize usage of high-capacity tapes such as the IBM Magstar and StorageTek 9840.

CONTINUED . . .

10.10 CONTINUED

**DUMP MULTIPLE
VOLUMES
CONCURRENTLY**

Backup 3 disk volumes concurrently to 3490E cartridge drives. The disks may be a combination of 3390s and 3380s (or any other supported disk type). For one volume (333333), a duplicate backup is created on another tape, so a total of 4 tape drives are required. If the tape drives are connected via ESCON or FICON channels, FDR software compression (COMPRESS=) is not recommended; however, tape hardware compression is recommended. Hardware compression is requested by the DCB=TRTCH=COMP parameter, but it may be the default in your installation.

```
//DUMP          EXEC   PGM=FDR,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=111111
//TAPE1         DD     UNIT=3490,DSN=BACKUP.V111111,DISP=(,CATLG),
//              VOL=(,,99),DCB=TRTCH=COMP
//SYSPRIN1      DD     SYSOUT=*
//DISK2         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=222222
//TAPE2         DD     UNIT=3490,DSN=BACKUP.V222222,DISP=(,CATLG),
//              VOL=(,,99),DCB=TRTCH=COMP
//SYSPRIN2      DD     SYSOUT=*
//DISK3         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=333333
//TAPE3         DD     UNIT=3490,DSN=BACKUP.V333333.COPY1,DISP=(,CATLG),
//              VOL=(,,99),DCB=TRTCH=COMP
//TAPE33        DD     UNIT=3490,DSN=BACKUP.V333333.COPY2,DISP=(,CATLG),
//              VOL=(,,99),DCB=TRTCH=COMP
//SYSPRIN3      DD     SYSOUT=*
//SYSIN         DD     *
                DUMP    TYPE=FDR,DSNENQ=USE,MAXTASKS=3
```

**DUMP
MULTIPLE
VOLUMES
CONCURREN-
TLY WITH
EXHPDM**

Backup 3 disk volumes concurrently using the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM is invoked by the SUBSYS= operands on the TAPEX DD statements; see [Section 80.33](#) and the ExHPDM program documentation for details on the values to provide. These 3 concurrent backups will be combined into one file on a tape managed by ExHPDM.

```
//DUMP          EXEC   PGM=FDR,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=111111
//TAPE1         DD     DSN=BACKUP.V111111,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN1      DD     SYSOUT=*
//DISK2         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=222222
//TAPE2         DD     DSN=BACKUP.V222222,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN2      DD     SYSOUT=*
//DISK3         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=333333
//TAPE3         DD     DSN=BACKUP.V333333.COPY1,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN3      DD     SYSOUT=*
//SYSIN         DD     *
                DUMP    TYPE=FDR,DSNENQ=USE,MAXTASKS=3,MAXERR=1
```

10.10 CONTINUED

**DUMP ONE
VOLUME WITH
TAPE
HARDWARE
COMPRESSION**

Dump a 3390 disk volume to a 3590 (Magstar) tape cartridge using the 3590 hardware compression feature. The tape will be retained by tape management software for 31 days. Since 3590s always attach via ESCON or FICON channels, you always want to use hardware compression instead of FDR software compression (COMPRESS=). Hardware compression is the default on 3590s, but DCB=TRTCH=COMP is specified to be sure it is used.

```
//DUMP          EXEC   PGM=FDR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=3390,DISP=OLD,VOL=SER=ABR123
//TAPE1         DD     UNIT=3590-1,DSN=BACKUP.V123456,DISP=(,CATLG),
//              DCB=TRTCH=COMP,LABEL=RETPD=31
//SYSIN         DD     *
DUMP           TYPE=FDR
```

**DUMP WITH
RESERVE**

Dump four disk volumes to two tape drives concurrently. Since the same tape volume and drive are referenced by TAPE1 and TAPE2 (VOL=REF=), the disks DISK1 and DISK2 will be backed up as separate files on that tape, one at a time. TAPEA and TAPEB refer to the same tape drive (UNIT=AFF=) but each backup will request a new scratch tape and DISKA and DISKB will be backed up, one at a time. But the ATTACH option allows a backup on one tape drive to execute concurrently with the backup on the other drive, using internal subtasking. The disk volume VTOCs will be ENQed and the disk devices reserved during the backup. The number of errors is set to 1, so that any volume that encounters any I/O error will terminate immediately.

```
//DUMP          EXEC   PGM=FDR,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=SYSALLDA,DISP=OLD,VOL=SER=SYSRES
//TAPE1         DD     UNIT=TAPE,DISP=(,CATLG),DSN=BACKUP.SYSRES
//SYSPRIN1      DD     SYSOUT=*
//DISK2         DD     UNIT=SYSALLDA,DISP=OLD,VOL=SER=PAGE01
//TAPE2         DD     VOL=REF=*.TAPE1,DISP=(,CATLG),DSN=BACKUP.PAGE01,
//              LABEL=2
//SYSPRIN2      DD     SYSOUT=*
//DISKA         DD     UNIT=SYSALLDA,DISP=OLD,VOL=SER=PACK01
//TAPEA         DD     UNIT=TAPE,DISP=(,CATLG),DSN=BACKUP.PACK01
//SYSPRINA      DD     SYSOUT=*
//DISKB         DD     UNIT=SYSALLDA,DISP=OLD,VOL=SER=MVS002
//TAPEB         DD     UNIT=AFF=TAPEA,DISP=(,CATLG),DSN=BACKUP.MVS002,
//              LABEL=2
//SYSPRINB      DD     SYSOUT=*
//SYSIN         DD     *
DUMP           TYPE=FDR,ENQ=RESERVE,MAXERR=1,ATTACH
```

**COPY FDR
BACKUPS**

Copy a FDR backup data set, creating an additional copy. **FDR backups cannot be copied by standard copy utilities, such as IEBGENER; such a copy may appear to complete successfully but the copy will be unusable.** FDRTCOPY is the tape copy utility provided with FDR. Complete documentation on FDRTCOPY is in [Section 60](#) of this manual. In this example, the input backup (TAPEIN) could be any FDR or DSF backup on tape or disk and the output (TAPEOUT) is on tape. You might want to use FDRTCOPY instead of using the duplication backup option of FDR to reduce the number of tape drives required by FDR backups. You also might use FDRTCOPY to create a third copy (or more).

```
//FDRTCOPY      EXEC   PGM=FDRTCOPY,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//TAPEIN        DD     DSN=BACKUP.VMVS001.COPY1,DISP=OLD
//TAPEOUT       DD     DSN=BACKUP.VMVS001.COPY2,UNIT=TAPE,
//              VOL=(,,20),DISP=(,CATLG)
//SYSIN         DD     *
COPY
```

CONTINUED ...

10.11 FDR RESTORE EXAMPLES

**FULL VOLUME
RESTORE
CHANGING
THE VOLUME
SERIAL**

Restore a 3390 and a 3380 volume, changing the existing volsers of the output disk volumes (DATA01 and DATA80) to the volsers of the volumes being restored (DATA12 and DATA85); if the original volume serials are still online, these output disks will be placed offline at the completion of the restore. The backups are not cataloged, so the backup tape volume serials are provided on the TAPEX DD statements. FDR will issue a WTOR to the MVS operator (message FDRW01) requesting permission to perform the restore.

```
//RESTORE      EXEC   PGM=FDR,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=3390,DISP=OLD,VOL=SER=DATA01
//TAPE1        DD     UNIT=TAPE,DSN=BACKUP.VDATA12,DISP=OLD,
//              VOL=SER=(T00100,T00110,T00120,T00130)
//DISKA        DD     UNIT=3380,DISP=OLD,VOL=SER=DATA80
//TAPEA        DD     UNIT=TAPE,DSN=BACKUP.VDATA85,DISP=OLD,
//              VOL=SER=(T00200,T00205,T00235)
//SYSIN        DD     *
RESTORE        TYPE=FDR,CPYVOLID=YES
```

**FULL VOLUME
RESTORE
RETAINING
THE OUTPUT
SERIAL**

Restore a disk from a tape backup without changing the volume serial of the output disk volume. The volume will be relabeled later to the volume serial of the volume which was dumped so VOLRESET=NO specifies that FDR will not rename the VTOCIX or ABR model. Neither the backed-up disk nor the target disk can be SMS-managed. FDR will issue a WTOR to the MVS operator (message FDRW01) requesting permission to perform the restore.

```
//RESTORE      EXEC   PGM=FDR,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=SYSDA,DISP=OLD,VOL=SER=ALTRES
//TAPE1        DD     DSN=BACKUP.VMVSRES,DISP=OLD
//SYSIN        DD     *
RESTORE        TYPE=FDR,CPYVOLID=NO,VOLRESET=NO
```

**RESTORE SMS-
MANAGED
VOLUME**

Restore a SMS-managed disk from a backup tape. Since the backup tape indicates it is a backup of a SMS-managed volume, the output volume must be marked as SMS-managed as well. Since a full-volume restore of an SMS volume to a new volser would result in uncataloged data sets in violation of SMS rules, CPYVOLID=YES is forced. FDR will not ask for operator permission; the restore will begin immediately.

```
//RESTSMS      EXEC   PGM=FDR,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=DISK,DISP=OLD,VOL=SER=SMS001
//TAPE1        DD     DSN=BACKUP.VSMS001,DISP=OLD
//SYSIN        DD     *
RESTORE        TYPE=FDR,CONFMESS=NO
```

10.11 CONTINUED

**RESTORE TO
LARGER DISK**

Convert a disk volume from a 3390-2 (2226 cylinders) to a 3390-3 (3339 cylinders), by restoring from a backup of the 3390-2. The current volume serial of the 3390-3 (D3390A) will be changed to the volume serial of the 3390-2 (D3390B). If the original volume is still mounted at the time of the restore, then the new volume will automatically be placed offline at the end of the restore; at that time, the original volume should be VARYed offline, and the new volume should be MOUNTed. Free space will automatically be adjusted to include the additional cylinders. If the original volume had an indexed VTOC, it must be rebuilt on the new volume (see next example).

```
//RESTORE      EXEC  PGM=FDR,REGION=1M
//SYSPRINT     DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//DISK1        DD    UNIT=3390,DISP=OLD,VOL=SER=D3390A
//TAPE1        DD    DSN=BACKUP.D3390B,DISP=OLD
//SYSIN        DD    *
RESTORE        TYPE=FDR,CPYVOLID=YES
```

**CONCURRENT
RESTORES**

Restore 3 disk volumes concurrently. The 3 backups will be mounted on 3 separate tape drives and will be restored in parallel.

```
//RESTORE      EXEC  PGM=FDR,REGION=1M
//SYSPRINT     DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//DISK1        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA01
//TAPE1        DD    DSN=BACKUP.VDATA01,DISP=OLD
//SYSPRIN1     DD    SYSOUT=*
//DISKA        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA85
//TAPEA        DD    DSN=BACKUP.VDATA85,DISP=OLD
//SYSPRINA     DD    SYSOUT=*
//DISKY        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA89
//TAPEY        DD    DSN=BACKUP.VDATA89,DISP=OLD
//SYSPRINY     DD    SYSOUT=*
//SYSIN        DD    *
RESTORE        TYPE=FDR,MAXTASKS=3
```

**CONCURRENT
RESTORES
WITH EXHPDM**

Restore 3 disk volumes concurrently, where the backups were done with the ExHPDM software product from StorageTek (see example in [Section 10.11](#)). ExHPDM will mount the tape or tapes required to restore these backups. The restore will be most efficient if all the backups which were directed to a single tape by ExHPDM are restored concurrently, but it is possible to restore one or more of them from that tape.

```
//RESTORE      EXEC  PGM=FDR,REGION=1M
//SYSPRINT     DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//DISK1        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA01
//TAPE1        DD    DSN=BACKUP.VDATA01,DISP=OLD,SUBSYS=SOV
//SYSPRIN1     DD    SYSOUT=*
//DISKA        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA85
//TAPEA        DD    DSN=BACKUP.VDATA85,DISP=OLD,SUBSYS=SOV
//SYSPRINA     DD    SYSOUT=*
//DISKY        DD    UNIT=DISK,DISP=OLD,VOL=SER=DATA89
//TAPEY        DD    DSN=BACKUP.VDATA89,DISP=OLD,SUBSYS=SOV
//SYSPRINY     DD    SYSOUT=*
//SYSIN        DD    *
RESTORE        TYPE=FDR,MAXTASKS=3,MAXERR=1
```

CONTINUED ...

10.11 CONTINUED

**REBUILD
INDEXED
VTOC**

Rebuild the indexed VTOC (VTOCIX) on any disk. This will be required after any FDR restore or copy a disk to another disk with a different capacity (number of cylinders). It will also be required after COMPAKTion to a new volume with CPYVOLID=YES, if the original volume serial was still online at the time. This example assumes that the SYS1.VTOCIX.vvvvvv data set exists on the volume but has been disabled.

```
//BUILDIX      EXEC   PGM=ICKDSF
//SYSPRINT     DD     SYSOUT=*
//DISK1        DD     UNIT=SYSALLDA,DISP=OLD,VOL=SER=PROD01
//SYSIN        DD     *
    BUILDIX    DDNAME(DISK1)  IXVTOC
```

10.12 FDR DISK COPY EXAMPLES

**COPY A DISK
VOLUME WITH
DUPLICATE
BACKUP**

Copy a 3390 volume to another 3390 volume, creating a backup on tape at the same time. CPYVOLID=NO and VOLRESET=YES are specified (but they are also the defaults) so the volume serial of the disk being copied to will be retained and the VTOCIX and ABR Model DSCB will be renamed to match the new volser. This is appropriate when you do not plan to eventually relabel the output volume back to the volser of the input, **however, VSAM clusters will become unusable**. All of the data sets on the volume being copied will be enqueued during the copy. Neither volume can be SMS-managed.

```
//COPY          EXEC  PGM=FDR,REGION=1M
//SYSPRINT      DD    SYSOUT=*
//SYSPRINT1     DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    UNIT=3390,DISP=OLD,VOL=SER=SYSRES
//TAPE1         DD    UNIT=3390,DISP=OLD,VOL=SER=DUPRES
//TAPE11        DD    UNIT=TAPE,DISP=(,CATLG),DSN=SYSRES.BACKUP,VOL=(,,20)
//SYSIN         DD    *
COPY           TYPE=FDR,DSNENQ=HAVE,CPYVOLID=NO,VOLRESET=YES
```

**COPY A DISK
VOLUME
CHANGING
THE VOLUME
SERIAL**

Copy a 3390 volume to another 3390 volume of the same size. The volume serial number of the receiving volume will be changed to the volume being copied. The output volume will automatically be set offline at the end of the restore. To replace the original with the copy, you must vary the original volume offline on all systems, and vary on the new copy.

```
//COPY          EXEC  PGM=FDR,REGION=1M
//SYSPRINT      DD    SYSOUT=*
//SYSPRINT1     DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    UNIT=3390,DISP=OLD,VOL=SER=DISK01
//TAPE1         DD    UNIT=3390,DISP=OLD,VOL=SER=DISK02
//SYSIN         DD    *
COPY           TYPE=FDR,CPYVOLID=YES
```

**COPY TO
LARGER DISK**

Convert a 3390-2 (2226 cylinders) to a 3390-3 (3339 cylinders) volume, with a direct disk-to-disk copy, creating a backup copy on 3490E cartridges as well. The volume serial of the 3390-3 (D3390A) will be changed to the volume serial of the 3390-2 (D3390B), and the new volume will automatically be placed offline at the end of the copy. At that time, the original volume should be VARYed offline, and the new volume should be MOUNTed. Free space will automatically be adjusted to include the additional cylinders. If the original volume had an indexed VTOC, it must be rebuilt on the new volume (See example in [Section 10.11](#)).

```
//COPY          EXEC  PGM=FDR,REGION=1M
//SYSPRINT      DD    SYSOUT=*
//SYSPRINT1     DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    UNIT=3390,DISP=OLD,VOL=SER=D3390B <--- 3390-2
//TAPE1         DD    UNIT=3380,DISP=OLD,VOL=SER=D3390A <--- 3390-3
//TAPE11        DD    UNIT=3490,DISP=(,CATLG),
//              DSN=BACKUP.D3390B,VOL=(,,20)
//SYSIN         DD    *
COPY           TYPE=FDR,CPYVOLID=YES
```

10.12 CONTINUED

**COPY TO
SMALLER DISK**

Convert a 3380-K (2655 cylinders) to a 3380-E (1770 cylinders) volume, with a direct disk-to-disk copy. FDR will not normally allow a copy or restore to a smaller density disk; COMPAKTOR is recommended for that function. However, if the larger disk has no data sets allocated past the end of the smaller disk (cylinder 1769 in this example), FDR can do this copy or restore with PROT=NONE.

All of the comments in the previous example (COPY TO LARGER DISK) apply.

```
//COPY          EXEC  PGM=FDR,REGION=1M
//SYSPRINT      DD    SYSOUT=*
//SYSPRIN1      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1        DD    UNIT=3380,DISP=OLD,VOL=SER=D3380B <--- 3380-K
//TAPE1         DD    UNIT=3380,DISP=OLD,VOL=SER=D3380A <--- 3380-E
//SYSIN        DD    *
                COPY   TYPE=FDR,CPYVOLID=YES,PROT=NONE
```

10.13 FDR INSTANTBACKUP EXAMPLES

This section contains some examples of the use of FDR with FDR InstantBackup. FDR InstantBackup is described in more detail in [Sections 25 through 29](#), depending on the type of disk subsystem you are using; there are more examples for each subsystem in those sections.

FDR InstantBackup is an additional cost enhancement to FDR.

INSTANT BACKUP OF VOLUME WITH SNAPSHOT

The disk to be backed up is in an IBM RVA or StorageTek Iceberg/SVA with the Snapshot feature. The step SNAP will snap PROD01 to 01FA, creating a point-in-time image of the volume. Step BACKUP will backup the snapped copy of volume PROD01. After the backup is complete, FDR will release all of the back-end (internal disk) storage associated with the snapped volume. [See Section 26](#) for details of FDR InstantBackup for RVA/Iceberg/SVA.

```
// SNAP      EXEC      PGM=FDRSNAP
//SYSPRINT DD      SYSOUT=*
//TAPE1      DD      DUMMY
//SYSIN      DD      *
      SNAP TYPE=FDR
      MOUNT VOL=PROD01, SNAPUNIT=01FA
//BACKUP     EXEC      PGM=FDR, REGION=0M, COND=(0, NE, SNAP)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1      DD      UNIT=SYSALLDA, VOL=SER=PROD01, DISP=OLD
//TAPE1      DD      DSN=BACKUP.VPROD01(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSIN      DD      *
      DUMP TYPE=FDR, SNAP=(USE, REL)
```

INSTANT BACKUP OF EMC SYMMETRIX

The disk to be backed up is in an EMC Symmetrix with the Timefinder feature. The BCV at address 01FA has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous one-time ESTABLISH has been issued to establish that pairing. The step SPLIT will split the BCV from its online volume and wait for the split to complete, creating a point-in-time image of the volume. Step BACKUP will backup the BCV copy of volume PROD01 and re-synchronize the BCV with the online volume. [See Section 25](#) for details of FDR InstantBackup for EMC Symmetrix.

```
//SPLIT      EXEC      PGM=EMCTF
//SYSOUT DD      SYSOUT=*
//SYSIN      DD      *
      SPLIT 1,01FA, WAIT
//BACKUP     EXEC      PGM=FDR, REGION=0M, COND=(0, NE, SPLIT)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1      DD      DSN=FDR.USE.UNIT01FA, UNIT=SYSALLDA,
//              VOL=SER=PROD01, DISP=OLD
//TAPE1      DD      DSN=BACKUP.VPROD01(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSIN      DD      *
      DUMP TYPE=FDR, BCV=(USE, RET)
```

Details of FDR InstantBackup for supported disk systems are found in [Sections 25-29](#).

10.13 CONTINUED

DUMP WITH HSDM COMPRESSION Dump a disk volume in a disk subsystem with the HSDM (High Speed Data Mover) option installed. FDR will backup internal compressed images of each data track, reducing elapsed time up to 60%.

```
//DUMP          EXEC  PGM=FDR,REGION=0M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1        DD    UNIT=3390,DISP=OLD,VOL=SER=ABR123
//TAPE1         DD    UNIT=TAPE,DSN=BACKUP.V123456,DISP=(,CATLG)
//SYSIN         DD    *
                DUMP   TYPE=FDR,DCT=YES
```

10.25 FDRQUERY - FDR/ABR STATISTICS QUERY

The FDR/ABR statistics query program (FDRQUERY) has been specifically designed to produce DUMP statistics as if a user were backing up disk volumes with FDR or with FDRABR using TYPE=ABR. Also you can request how many data sets would be archived by FDRABR using last reference date groupings.

**BACKUP
SIMULATION**

The FDR statistics query program will scan all of the volumes specified comparing the number of tracks which would be dumped if FDR were to execute against the volume as compared to ABR dumping only data sets which have the update indicator. The saving is printed in tracks and percentage. Since ABR is usually run every day, the query program will not report on any data set with an update indicator if it has not been referenced in the last two days. Even using this technique, the query program may indicate that a larger number of data sets will be dumped by ABR than would actually take place, if the update indicator is on but the data set was only read in the last two days.

**ARCHIVE
SIMULATION**

The FDR statistics query program can scan all of the volumes specified to report on the number of data sets and the tracks they occupy, grouped by the last time they were referenced. As a default the query program will report data sets in 30 day groups. The purpose of this report is to show an FDR user how much disk space could be saved, if ABR were used to archive off data sets which have not been referenced in a specific period of time.

**SAMPLE
ARCHIVE
REPORT****SUMMARY LEVEL -- VOLUME SERIAL NUMBER**

<u>VOLSER</u>	<u>DEVTYPE</u>	<u>ALLOC TRACKS</u>	<u>BEFORE %ALLOC</u>	<u>AFTER %ALLOC</u>	<u>LAST USED</u>		<u>SAVINGS IF ARCHIVED</u>		
					<u>DAYS</u>	<u>DATE</u>	<u>DSNS</u>	<u>TRACKS</u>	<u>%SAVED</u>
PROD32	3380-K	17013	42.71%	32.08%	30	98051	12	4237	24.90%

The report fields are:

ALLOC TRACKS -- This number represents the total currently allocated tracks.

BEFORE % ALLOC -- This number is the percentage of the volume that is in use. It is the number of tracks allocated, divided by the total number of tracks on the disk. For example, volume PROD32 (a 3380-K DASD) contains 39,825 tracks of which 17,013 tracks are allocated. The BEFORE % ALLOC is 42.71% ($17,013 \div 39,825 = 42.71\%$).

AFTER % ALLOC -- This field represents the percent allocated the volume will be if you archive data sets that have not been used for the number of days specified in the next column. For example, if on PROD32, there are 12 data sets with 4237 tracks that have not been referenced in 30 days and you archive them, the AFTER % ALLOC will be 32.08%.

$$\begin{array}{ccccccc}
 17,013 & \text{---} & 4237 & = & 12,776 & \div & 39,825 & = & 32.08\% \\
 \text{Allocated} & & \text{Archived} & & \text{Tracks} & & \text{Number} & & \text{After \%} \\
 \text{Tracks} & & \text{Tracks} & & \text{after} & & \text{of} & & \text{ALLOC} \\
 & & & & \text{Archiving} & & \text{Tracks} & &
 \end{array}$$

LAST USED -- The DAYS field shows the number of days used to calculate this line in the report. The data sets on this line have not been used for this number of days. The DATE field shows the corresponding Julian date (today's date minus DAYS).

SAVINGS IF ARCHIVED -- These fields show the number of data sets which would be archived, the number of tracks allocated to those data sets and the percentage of the total allocated tracks they represent.

For example, if volume PROD32 has 17,013 tracks allocated and you archive 4237 tracks, the % SAVED will be 24.90% ($4237 \div 17,013 = 24.90\%$). This means that 24.90% of the allocated space would be freed by ARCHIVING.

10.26 FDRQUERY JOB CONTROL

EXEC STATEMENT	Must specify the name of the FDR/ABR DUMP statistics query program (PGM=FDRQUERY) and may also contain a region requirement of 512K.
STEPLIB or JOB LIB DD STATEMENT	If FDR is not in the system linklist, specifies the FDR program library. This should be an APF authorized library.
SYS PRINT DD STATEMENT	Specifies the primary output message data set. This is a required DD statement and is usually a SYSOUT data set.
SYS IN DD STATEMENT	Specifies the control statement data set. Usually an input stream or DD * data set.

10.27 FDRQUERY REPORT STATEMENT

REPORT	ARCHIVE BACKUP	,MAXONLINE=nnnn
	,AGE=nnn	,ONLINE
	,AGEINC=nnn	,STORGRP=CCCCCCCC
	,LINECNT=nn	,UNITNAME=CCCCCCCC
	,LRDAYS=nnn	,VOL=vvvvvv VOLG=vvvvv

**REPORT
STATEMENT**

The REPORT statement of FDRQUERY is used to print total tracks and data sets by disk volume and device, simulating the FDR/FDRDSF/FDRABR DUMP and ARCHIVE Commands.

FDRQUERY will default to scanning all of the disk volumes which are currently online (unless overridden by VOL(G), STORGRP, or UNITNAME). The report will be formatted for an 80 byte terminal screen.

OPERANDS ARCHIVE

FDR will summarize Archive statistics by disk volume, indicating the number of data sets and tracks that would be freed up if FDRABR archive was executed against these volumes. The data sets are grouped by last reference date, according to the values of AGE= and AGEINC=.

BACKUP

FDRQUERY will print summarized DUMP statistics by disk volume and device type, indicating the number of tracks and data sets that would be dumped by FDR and comparing the results to FDRABR dumping only the data sets which have been updated. The used portion of a data set (partitioned or sequential) is the number of tracks encompassed by the Last Block Pointer (DS1LSTAR). Data sets are considered to be 'UPDATED' if the UPDATE flag is on in the Format 1 DSCB and the data set has been referenced within the number of days specified by LRDAYS=.

NOTE: The 'BACKUP' operand conflicts with the operand 'ARCHIVE'. Either the operand BACKUP or ARCHIVE must be specified.

AGE=

Specifies the starting number of days since a data set has been referenced as used by the AGING summary in the archive simulation. The number may be from 1 to 999 inclusive.

Default is 30.

AGEINC=

Specifies the number to be added to the age value to derive the next date control break in the AGING summary in the archive simulation. The number may be from 1 to 999 inclusive.

Default is 30.

NOTE: AGE and AGEINC operands are ignored if 'BACKUP' is specified.

LINECNT=

Specifies the maximum number of lines to be printed on any report page. The number may be from 28 to 99 inclusive.

Default is 58.

10.27 CONTINUED

LRDAYS=	<p>BACKUP only. Specifies the number of days that a data set must have been referenced for the update indicator to be considered for the ABR portion of the report. May specify from 0 to 999 days, with zero meaning today.</p> <p>Default is 1, the data set must have been referenced in the last two days.</p>
MAXONLINE=	<p>Specifies the maximum number of disk volumes that may be allocated via DD statements and/or by dynamic allocation during any single program execution. The number may be from 5 to 9000 inclusive.</p> <p>The default is 256.</p>
ONLINE	<p>ONLINE specifies that the query program is to use all available ONLINE DASD devices to satisfy the user-specified selection criteria.</p> <p>Default: ONLINE unless you specify one or more of the operands VOL=, VOLG=, STORGRP= and/or UNITNAME=.</p>
STORGRP=	<p>Specifies that the disk volumes that are part of the SMS storage group name specified are to be processed.</p>
UNITNAME=	<p>Specifies that all disk volumes mounted on a device which is defined to be part of the esoteric or generic unit name specified are to be processed, e.g., UNITNAME=SYSDA or UNITNAME=3390.</p> <p>NOTE: Selection by UNITNAME requires that FDRQUERY be installed in an authorized library with an authorization code of one (1). If UNITNAME selection is attempted and FDRQUERY is not authorized, an FDR640 message will be issued and the request bypassed.</p>
VOL=	<p>Specifies the volume serial of the disk to be processed.</p>
VOLG=	<p>Specifies the prefix of disk volumes serials to be processed; any online volume whose serial starts with the prefix will be included. Multiple prefixes may be specified by enclosing them in parenthesis separated by commas.</p>

10.28 FDRQUERY EXAMPLES

**ALL ONLINE
DISKS**

Print DUMP and ARCHIVE statistics for all ONLINE disks.

```
//STEP1      EXEC   PGM=FDRQUERY
//SYSPRINT    DD     SYSOUT=A
//SYSIN       DD     *
REPORT       BACKUP
REPORT       ARCHIVE
```

**SPECIFIC DISK
VOLUMES**

Print comparison DUMP statistics for specific disk volumes.

```
//STEP1      EXEC   PGM=FDRQUERY
//SYSPRINT    DD     SYSOUT=A
//SYSIN       DD     *
REPORT       BACKUP,VOL=(PRODPK,LIB501,TEST01)
```

**VOLSER
PREFIX**

Print ARCHIVE statistics for all ONLINE disk volumes starting with the prefix 'PROD'. Increment the ADAYS/ADATE range by 15 days.

```
//STEP1      EXEC   PGM=FDRQUERY
//SYSPRINT    DD     SYSOUT=A
//SYSIN       DD     *
REPORT       ARCHIVE,VOLG=PROD,AGEINC=15
```

**SMS STORAGE
GROUP**

Print ARCHIVE statistics for all disk volumes within the SMS storage group TSOTEST. Set the minimum number of days since last referenced to 60.

```
//STEP1      EXEC   PGM=FDRQUERY
//SYSPRINT    DD     SYSOUT=A
//SYSIN       DD     *
REPORT       ARCHIVE,AGE=60,STORGRP=TSOTEST
```

UNIT NAME

Print comparison DUMP statistics for all disk volumes within the esoteric unit named DISK.

NOTE: UNITNAME selection requires APF authorization.

```
//STEP1      EXEC   PGM=FDRQUERY
//SYSPRINT    DD     SYSOUT=A
//SYSIN       DD     *
REPORT       BACKUP,UNITNAME=DISK
```

**INVOKE FROM
TSO**

Report to a TSO terminal the archiving or DUMP statistics for all online volumes. The report will be formatted for an 80 byte screen. The SYSIN and SYSPRINT data sets must be allocated to the terminal prior to the CALL (EXAMPLE: ALLOCATE DD(SYSIN) DA(*)).

```
CALL 'FDR library name (FDRQUERY)' "PRESS" "ENTER"
FDRQUERY--ENTER COMMAND OR END
REPORT BACKUP "PRESS" "ENTER"
OR
REPORT ARCHIVE "PRESS" "ENTER"
```

These commands will report on all online volumes. Use VOL= or VOLG= to limit selection criteria (i.e. REPORT ARCHIVE,VOLG=TSO)

10.29 FDRQUERY SUMMARY ILLUSTRATION

FDRQUERY INCREMENTAL BACKUP SAVINGS

VOLSER	DEVTYPE	TRACKS	ALLOCATED			INCREMENTAL			SAVINGS		
			TRACKS	DSNS	PER	TRACKS	DSNS	PER	TRACKS	DSNS	PER
MVSYS1	3380	13275	7359	124	56%	3947	91	30%	3412	33	46%
MVSYS2	3380	13275	11223	97	85%	1634	18	12%	9586	79	86%
TSO001	3350	16650	16125	968	96%	4015	29	24%	12110	939	75%
TSO002	3350	16650	15500	812	93%	3900	150	23%	11600	662	74%

Tape and elapsed time savings over full volume dump

SUMMARY LEVEL BY DEVICE TYPE

VOLSER	DEVTYPE	TOTAL TRACKS	ALLOCATED			INCREMENTAL			SAVINGS		
			TRACKS	DSNS	PER	TRACKS	DSNS	PER	TRACKS	DSNS	PER
TOTAL	3350	33300	31625	1780	95%	7915	179	24%	23710	1601	75%
TOTAL	3380	26550	18582	221	70%	5581	109	21%	13001	112	70%

ABR's Incremental Backup & Recovery

ABR automates the backup of disk volumes. Data sets are automatically backed up when updated. ABR incrementals can save you 50-80% of your backup time compared to DFHSM.

Data Set Restore – ABR can automatically recover data set(s) from the most current or older backups. Entire disk volumes can be recreated as if a full volume backup was taken the prior night.

FDRQUERY SPACE MANAGEMENT SAVINGS

VOLSER	DEVTYPE	ALLOC TRACKS	BEFORE %ALLOC	AFTER %ALLOC	LAST USED		SAVINGS IF ARCHIVED		
					DAYS	DATE	DSNS	TRACKS	%SAVED
TSO002	3380-K	36045	90.51%	43.19%	30	89051	6901	18841	52.27%
				54.94%	60	89021	4787	14162	39.29%
				63.71%	90	88356	2962	10672	29.61%
				68.58%	120	88326	1897	8730	24.22%

PROD32	3380-K	27479	69.00%	49.01%	30	89051	51	7959	40.78%
				32.08%	60	89021	12	4237	24.90%

Space occupied by data sets inactive for 30 days.

SUMMARY LEVEL BY DEVICE TYPE

VOLCNT	DEVTYPE	ALLOC TRACKS	BEFORE %ALLOC	AFTER %ALLOC	LAST USED		SAVINGS IF ARCHIVED		
					DAYS	DATE	DSNS	TRACKS	%SAVED
6	3380-K	173095	72.44%	52.66%	30	89051	9423	47264	27.30%
				55.81%	60	89021	5296	39722	22.94%
				60.48%	90	88356	4199	28558	16.49%
				64.94%	120	88326	2972	17911	10.34%

ABR's Archive & Auto Recall

Multi-Level Archive – ABR can simultaneously archive data set(s) to disk and tape. The data is stored in backup format, which generally requires significantly less space on disk than originally used. The data is retained on disk for a short

user-specified time period, and is then automatically deleted by ABR, leaving the tape copy for a longer retention period.

Auto Recall -- ABR will automatically recall ARCHIVED data sets when they are referenced by a TSO or BATCH user.

15.01 SAR TECHNICAL DESCRIPTION**SAR
OVERVIEW**

SAR is the Stand Alone portion of the FDR DASD Management System. SAR is an IPLable, self-loading, stand-alone program not requiring an operating system.

SAR enables you to BACKUP or RESTORE a disk volume when, due to the unavailability of an operating system, FDR cannot be run. SAR may be used to recover or move disk volumes after DASD failures or HDA replacements. It may also be used at Disaster Recovery sites, or at new primary sites, to restore the disk volumes necessary to IPL.

SAR is simple to use and provides console messages to prompt you, so there are no control cards to prepare. Diagnostic messages are issued if SAR encounters difficulty during the BACKUP or RESTORE.

SAR supports three major operations:

RESTORE **a disk volume from a backup tape.**

DUMP **an entire disk volume or selected data sets to tape.**

CLIP **change the volume serial of a disk volume in place.**

**SUPPORTED
HARDWARE**

SAR will execute on any CPU capable of running any version of MVS or OS/390 and will operate on any type of disk supported by any of the supported MVS versions. SAR can operate in an LPAR or on a VM virtual machine.

SAR can work with any of these tape drives:

IBM 3420/3422 (9 track)

IBM 3490E (36 track cartridge)

StorageTek Timberline

StorageTek 9840

IBM 3480/3490 (18 track cartridge)

IBM 3590 (Magstar cartridge)

StorageTek Redwood

any tape drive which emulates one of the above.

On the IBM PC Server S/390 (also called a P/390), SAR supports the MVS disk volumes emulated on OS/2 disk storage. For tape, SAR can use either standard tape drives attached by a channel adaptor card, the 4MM tape drive built into the server (emulating a 3420 or 3480), or OS/2 disk files which are defined as emulating a tape. SAR also operates on a RS/6000 S/390 (also called a R/390) with similar functions, running under AIX. P/390 operations are described here; R/390 operations are not described but they are similar to the P/390.



From supporting SAR for the last 28 years, Innovation has a lot of experience on what problems can occur and how users react to try and solve them. We have put a lot of effort into making SAR easy to use. For example, SAR displays the available choices for most options, and prompts you if you enter an invalid value. Since SAR is used infrequently and since you are often rushed to complete SAR operation to recover from a system crash or DASD failure, proper procedures may not be followed. After reading the rest of the SAR documentation, please review [section 15.25](#) for suggestions on avoiding common problems.

15.01 CONTINUED

SAR RESTORE
OPERATIONS

SAR can restore:

- An entire disk volume from a full-volume backup tape created by FDR, ABR, or SAR.
- Ranges of specific tracks, from any type of backup tape created by FDR, DSF, ABR, or SAR.

If the input tape to a SAR restore is not a full-volume backup tape, SAR issues a warning message and you have the option to continue the restore. This usually means that the wrong tape was mounted.

SAR can only restore backups to the same disk type, e.g., 3380 to 3380, 3390 to 3390 (or to devices which emulate the same disk type, such as 3390 to RAMAC emulating 3390).

On a full-volume restore, you have 3 options controlling the volume serial of the disk volume after the restore:

- retain the original volume serial of the output disk volume
- restore the volume serial of the disk on the backup tape
- specify a new volume serial from the console

SAR can restore to an uninitialized volume, i.e., one that does not currently contain a volume label or VTOC.

If you restore an ABR-created full-volume backup tape using SAR, the ABR model DSCB on the volume will be downleveled. The Generation and Cycle will reflect the previous generation. Use the REMODEL command (see [Section 50.44](#)) to update the ABR model after you relPL MVS. You should not try to restore from an ABR-created incremental (data set) backup.

SAR can restore from any file on an ABR backup tape, using the file sequence number on the INPUT TAPE UNIT command. The file sequence number is shown in the FDRABRP PRINT CATLG report, and in the FDR305 message in the listing of the backup job (see "ABR Backup Tapes" in [Section 15.24](#))

SAR BACKUP
OPERATIONS

SAR can back up to tape:

- An entire disk volume. The backup will contain all allocated tracks as shown in the VTOC, including the VTOC and the label track (cylinder 0 track 0).
- Data sets, by dsname or group name (similar to an DSF or ABR data set backup).
- Ranges of specific tracks (similar to a backup created by DSF with FROM/TO).

SAR will request scratch tapes to be used for the backup. It will display the data set name, volume serial number and current expiration date of the backup tape mounted; you must manually insure that the tapes are in scratch status. You can specify the data set name to be placed in the tape labels.

The backups produced by SAR can be restored by SAR, FDR or ABR. Individual data sets can also be restored from this backup using DSF or ABR. However, VSAM files cannot be restored by cluster name.

NOTE: Although tapes created by SAR are logically equivalent to FDR tapes, there are formatting differences, so they cannot be successfully compared to tapes created with FDR or ABR.

CLIP
OPERATIONS

CLIP stands for Change Label In Place. A SAR CLIP will change the volume serial number in the label track (cylinder 0 track 0) of a disk volume. SAR will prompt you for the existing serial number and the new serial number. This procedure will not affect any other data on the volume. However, if the volume serial is changed on a volume containing ICF VSAM clusters or SMS-managed data sets, that data will not be accessible.

CLIP will work only on initialized volumes, i.e., those containing a volume label and VTOC.

CONTINUED . . .

15.01 CONTINUED

**ABSOLUTE
TRACK
OPERATIONS**

You can specify that SAR is to back up or restore specific tracks. From 1 to 10 ranges of tracks may be specified. SAR will prompt you for the ranges of track addresses to be processed.

Absolute track restore does not update the DSCBs in the VTOC, so you usually must restore ranges of tracks that represent data sets already in the VTOC. Care should be taken if the label track (cylinder 0 track 0) or a track in the VTOC is restored, since this may make the volume unusable.

**DATA SET
OPERATIONS**

You can specify that SAR is to back up specific data sets or groups of data sets. SAR will prompt you for the data set names. From 1 to 10 data sets or groups of data sets may be specified. A group prefix is specified by inserting an * at the end of the prefix. For example, specifying SYS1.* will instruct SAR to dump all data sets starting with 'SYS1.'. If the VTOC is to be dumped, specify ***VTOC; this will also dump cylinder 0 track 0 (the volume label track).

NOTE: SAR RESTORE does not support data set operations. Backups created by SAR data set dump can only be restored with FDRDSF or by SAR with the absolute track option (and only to data sets preallocated at the proper track locations, since SAR absolute track restore will not update the VTOC).

VM SUPPORT

SAR can be IPLed on a VM virtual machine. See the member \$\$SAR in the FDR ICL (Installation Control Library) for details on VM usage.

DISK SUPPORT

SAR will automatically determine the type and model of disk being dumped or restored. For emulated disks (such as IBM RAMAC, IBM RVA, EMC, and StorageTek disks), SAR will identify the model that they emulate.

SAR cannot restore to a disk with a different geometry (tracks/cylinder and track capacity), such as 3380 to 3390, but it is able to restore to a larger capacity disk with the same device geometry (same disk type), for example, from a backup of a 3390-2 to a 3390-3. SAR will automatically recognize whether the output volume being restored is the same capacity as the volume that was backed up. If the volume being restored to is larger, SAR will turn on the DOS bit, so that the first time a data set is allocated on that volume after the system is IPLed, the system will call the VTOC conversion routine, which will rebuild the VTOC to reflect the larger amount of available space. If the volume contained an indexed VTOC, the index will be disabled and you must use ICKDSF to rebuild the index after IPL.

Normally SAR will not allow you to restore a backup to a model with a smaller capacity than the disk that was backed up. However, as long as there were no data sets allocated on the original disk beyond the end of the target disk, you can force SAR to restore a higher capacity disk to a lower capacity disk of the same type. You must override the device type that SAR has correctly determined and substitute the type that was backed up. For example, if your backup was created from a 3390-3 and you want to restore to a 3390-2, then specify OUTPUT DISK DEVICE=3390-3.

Warning: Restore to a smaller disk will only work if all of the allocated tracks on the original disk are contained within the physical capacity of the output disk. Otherwise, SAR will give various error messages. Innovation recommends that this type of restore be done only if no other options are available.

SAR can restore backups created with the HSDM (High Speed Data Mover) hardware option available on some disk subsystems. [See Section 80.33](#) for details on HSDM.

CONTINUED . . .

15.01 CONTINUED

TAPE SUPPORT

SAR RESTORE automatically supports tapes that were compressed by the software data compression feature of FDR (COMPRESS=). There are no special considerations. SAR cannot create a tape with FDR compression, but it can read and create cartridges that are compressed by the tape hardware (IDRC).

For restore, SAR will automatically request the second and subsequent tape volumes if more than one volume was used on the backup; however, the operator must know what tape volumes are required for the restore of a particular disk volume and mount them in the proper order. During backup, SAR will request a new scratch output tape be mounted if the current volume is full. In either case, a message will be issued to the console requesting the next volume. SAR will automatically recognize when the new volume has been mounted and continue with the operation.

It is the operator's responsibility to ensure that the required input or output tapes are mounted in the correct order, and to record the volume serials of output tapes. If you are restoring from ABR full-volume backups, an FDRABRP "PRINT CATLG" report can be used to identify the tapes required for the restore.

SAR can use 1 or 2 tape drives for the Dump or Restore operation. To request 2 tape drives, specify "2" after the TAPE DEVICE parameter (e.g., INPUT TAPE DEVICE=3480,2). Allowing SAR to use 2 tape drives will reduce wall clock time by making it possible for the operator to premount the next tape. SAR will automatically start writing or reading the next tape while the prior one is being rewound and unloaded.

If your cartridge tape drive has an automatic cartridge loader (ACL) you can load the tapes required into the ACL before the SAR operation is started. The ACL cannot be in system mode. If it is in automatic mode, the next tape is automatically loaded when the previous tape is unloaded by SAR. In manual mode, the operator must press START to load each cartridge.

If the required tape volumes are in an Automated Tape Library (ATL) such as the IBM 3494/3495 and StorageTek silos, special procedures may be required to mount the tapes required for the backup or restore. The SAR IPL tape can also be in the ATL. Please consult each vendor's documentation for details:

- The IBM 3494 supports a "stand-alone mode" for a given drive, selected from the library console. In that mode you can tell it which tape volumes to mount, and it will fetch and mount them automatically. You can also insert tapes from outside the library (such as the SAR IPL tape).
- The IBM 3495 and StorageTek silos must be placed in manual mode, allowing you to enter the enclosure and manually mount tapes. If the required tapes are in the library, you must locate them and manually mount them. The 3495 has a locate function on its "manual mode" console, while StorageTek has an IPLable utility called "POST VOLSER to Location" which identifies tape locations.

SAR supports both SL (standard label) and NL (non-labeled) tapes. Since most tapes are labeled, documentation on the use of NL tapes is only in the member \$\$\$SAR in the FDR ICL (Installation Control Library).

SAR cannot restore from backups created with the ExHPDM software from StorageTek ([see Section 80.33](#)).

CONTINUED . . .

15.01 CONTINUED

**CONSOLE
SUPPORT**

Any 3270-compatible display connected to a channel-attached non-SNA 3x74 control unit or integrated console adapter can be used as the SAR console. Any standard MVS MCS console is supported by SAR.

On the IBM P/390, SAR will use the 3270 console emulated by Communication Manager/2.

Normally SAR will wait for some device to present an ATTENTION interrupt and then will attempt to use that device as a console. ATTENTION will be generated when you press ENTER on a display console, so you simply need to press ENTER on your console to identify it to SAR.

However, if you have other display terminals (such as TSO terminals) attached by non-SNA 3x74 control units, it is possible that another user will press ENTER before you have a chance to, and will suddenly get the SAR menu on his screen. This is often a problem with anxious users waiting for the system to come back up. There are other devices (such as CTC adaptors and 37x5 communication controllers) which also present ATTENTION and may confuse SAR. If you are unable to prevent such false interrupts, you can customize SAR with FDRSARLR, specifying the addresses of your consoles (up to 5) so that SAR will attempt to use one of them without waiting for an ATTENTION.

SAR supports printer/keyboard consoles, which display one line at a time (see the member \$\$SAR in the FDR ICL (Installation Control Library) for details). In addition, on most S/390 systems, SAR supports the limited-function operator message facility on the hardware console, known as the "SCLP Console" as a printer/keyboard-type console. See "SCLP Console" in [Section 15.02](#) for instructions.

**ERROR
RECOVERY**

Some errors cause immediate termination of the current SAR function, but SAR is able to continue after certain types of errors. If an error occurs, SAR will display a message detailing the error. If SAR is able to continue, the message will include the text "REPLY CONTINUE OR TERM" and you will be prompted for a response. If you reply CONTINUE (or just press ENTER since CONTINUE is the default), SAR will continue; however, depending on the error, data may be lost (not dumped or restored). If you reply TERM, SAR will terminate its current operation, enabling you to restart SAR for another function. If you have a hardcopy device assigned, and you want a dump of SAR's memory for diagnostic purposes, you can use the PSW RESTART hardware function ([see Section 15.03](#)) to obtain such a dump before replying.

CONTINUED . . .

15.02 SAR OPERATION REQUIREMENTS

Naturally, SAR requires a processor (CPU, or CPC, or LPAR) on which to operate. That processor must be idle (ready to IPL).

SAR requires these I/O devices: the IPL device, an operator console, a hardcopy device, a tape drive and a disk volume. All must be physically addressable by the processor, enabled and ready. Tapes must be offline to all other CPUs. If no hardcopy device is available, SAR can run without one, but if problems occur, diagnosis may be more difficult.

NOTE: If SAR is IPLed on a multi-processor CPU, SAR will actually execute only on the processor it was IPLed on. If SAR is executed on an LPAR, all these devices must be accessible from that LPAR.

**IPL
REQUIREMENTS**

SAR can be IPLed from a tape, disk or card reader (including VM virtual card readers).

A copy of SAR is always the first file on any FDR distribution tape so you may IPL from it, but it is a time-consuming procedure since it is a labeled tape. For efficiency and convenience, the SAR program should be copied to a disk or unlabeled tape using the program FDRSARLR ([Section 15.20](#)). IEBGENER can also be used to copy the SAR program file to an unlabeled tape.

On a P/390, SAR can be IPLed from tape drives attached by the channel adaptor card, from the built-in 4mm tape emulating a 3420 or 3480, or from a MVS disk volume emulated by OS/2 (FDRSARLR must be used to place SAR on that volume, of course). Also, OS/2 files can be configured to emulate a tape drive; you can copy SAR to such a file, then IPL from it as if it were a tape.

**CONSOLE
REQUIREMENTS**

By default, SAR attempts to use as the operator console the first device that presents an ATTENTION interrupt after IPL is completed. Normally, you do this by simply pressing ENTER on the desired console. On an P/390, activate the first CM/2-emulated 3270 session, press the right-hand mouse key, and select "ATTN" from the menu that is displayed. For the SCLP console, see the instructions below.

If you have other display terminals (such as TSO terminals) attached by non-SNA 3x74 control units, it is possible that another user will press ENTER before you do, and will suddenly get the SAR menu on his screen. There are other devices (such as CTC adaptors and 37x5 communication controllers) which also present ATTENTION and may confuse SAR. If you are unable to prevent such false interrupts by resetting these devices or putting them physically offline, you can preconfigure SAR with the addresses of your consoles (up to 5) with FDRSARLR and SAR will attempt to use one of them without waiting for an ATTENTION.

Once SAR identifies the console, the processing options and their defaults are displayed on the screen and the cursor is positioned to the first modifiable option. A default value may be accepted by pressing ENTER; the cursor will then be positioned to the next option line. If you need to override the defaults shown on a line, type over the data on that line and press ENTER. If the override is not accepted, a message will be issued and the cursor will be repositioned on the invalid entry. If the override is accepted, the cursor will be positioned to the next entry. Where applicable, SAR will display a list of acceptable values for the current option at the bottom of the screen.

SAR is not a true full-screen processor. You cannot use the cursor keys to move to a line other than the current one, enter multiple parameters without pressing ENTER, use the ERASE EOF key or go back to correct a previous parameter.

It is always necessary to press ENTER once for each option line, whether the defaults are accepted or overridden.

15.02 CONTINUED

**SCLP
CONSOLE**

In complex hardware configurations, especially Sysplex configurations, there may not be a 3270-type non-SNA console available for use with SAR on the IPLed CPU. In this case, there is a software console facility built into most S/390 systems which can be used to operate SAR; it is called the "SCLP console" function. However, it is awkward to use, especially on ES/9000 processors, so it should be used only when no other options are available. It is accessed through the hardware console, which is the service processor console on ES/9000s and the HMC (hardware management console) on 9672s.

Since these console functions are not full-screen 3270-type consoles, SAR treats them much like printer-keyboard consoles, displaying one message at a time and waiting for replies to be entered.

On a service processor console:

After IPLing SAR, you can switch to the SCLP console by keying:

F OPRMSG

on the service processor console.

To enter a reply, you must press the CANCEL key (ALT-PF1) and wait a few seconds for the cursor to move to the "SCP MESSAGE" line before entering your reply. Unlike a normal console, you cannot enter a null message by simply pressing ENTER; you must enter at least one character. So that you can accept the SAR default replies, SAR has been modified to treat a reply consisting of just one blank as a null reply, so you can accept the defaults by entering one blank space and ENTER.

There is no ATTENTION function. To cause SAR to recognize that you want to use the SCLP console, enter a null reply as shown above; the SAR messages should then start to appear. At the end of SAR operation, to restart SAR without reIPLing, use the same procedure.

On HMCs:

Double-click the icon for "Operating System Messages". This opens a window for the SCLP console function. If necessary, select the tab for the image on which SAR is running. This window has a SEND COMMAND button and a RESPOND button; you may click on either to send replies to SAR, but SEND COMMAND is preferred.

Note that trailing blanks are not sent to SAR. This is a problem mainly when you wish to enter 6 blanks for the VOLUME SERIAL= reply; to accommodate this, SAR will also accept a volser reply of an asterisk (*) as equivalent to 6 blanks. It is also a consideration when you are replying with a tape data set name; you cannot reply with a name shorter than the name found on the tape.

There is no ATTENTION function. To cause SAR to recognize that you want to use the SCLP console, enter a null message with SEND COMMAND; the SAR messages should then start to appear. At the end of SAR operation, to restart SAR without reIPLing, use the same procedure.

CONTINUED . . .

15.03 LOADING SAR

SAR IPL
PROCEDURES

To IPL SAR, the following procedures must be followed:

STEP 1

From Tape -- Mount and ready the SAR IPL tape on an available tape drive. Rewind it if necessary.

From Disk -- Device must be on-line and ready.

STEP 2

IPL SAR from the selected device. The procedure will vary depending on the type of your CPU, but the procedure will be similar to the procedure you use to IPL your normal operating system, except that the IPL device address will be the device selected in Step 1. Note that a SYSTEM RESET is not required before IPLing (IPL does an automatic SYSTEM RESET) and that either an IPL NORMAL or an IPL CLEAR may be used (IPL CLEAR is recommended).

If you are not familiar with the IPL procedure for your CPU, details can be found in the operator's documentation from the hardware vendor. However, for quick reference, here are brief descriptions for common IBM CPUs:

On a system with a separate hardware service console:

-- If necessary, select the LPAR (Logical partition) on which you wish to IPL SAR. SAR will only be able to access devices configured to this LPAR.

-- Enter the Service Language command (from any screen):

LOAD CPn address

where "CPn" is the target processor, and "address" is the device address of the SAR IPL device.

-- If the IPL is successful, you must move to a software console to operate SAR.

On a system with a HMC (Hardware Management Console):

-- To IPL, drag the icon for the CPC image on which to wish to run SAR to the "LOAD" task icon. You will be prompted for the IPL device address; enter the address where SAR resides.

On a P/390:

-- Make sure that CM/2 (Communication Manager/2) is started to provide console support.

-- If you have copied SAR to an OS/2 file, you may need to use the "P/390 Configuration" dialog to assign that file to a P/390 address as a tape drive.

-- Open an OS/2 window and issue the command "IPL address". You may need to switch to the P390 subdirectory to use the IPL command.

NOTE: If SAR is being IPLed from an INNOVATION FDR distribution tape, you must perform the IPL procedure 5 times. The first 4 IPLs will fail as it reads the header labels (VOL1, HDR1, HDR2, and tape mark); the fifth IPL should succeed (allow a few seconds after each error to be sure the processor has completed the IPL procedure before initiating the next one). We recommend copying SAR to a disk or unlabeled tape to avoid this.

15.03 CONTINUED

SAR IPL
PROCEDURES
(continued)

STEP 3

If the CONSOLE UNIT was specified via FDRSARLR ([Section 15.22](#)), the operator messages will automatically appear on the unit specified, if it is available and if it responds to console commands. If more than one console unit address was specified, each will be tried in turn. If none of the specified consoles respond, SAR acts as if no console unit was specified.

--OR ELSE--

When the IPL successfully completes, the processor should enter a wait state with a code of X 'FFFF' in the last two bytes of the PSW indicating SAR is waiting to identify a console (depending on your CPU type, you may not be able to see this PSW unless you STOP the CPU). At this time you can select the SAR console by pressing the ENTER key on a display console.

On an P/390, activate the first CM/2-emulated 3270 session, press the right-hand mouse key, and select "ATTN" from the menu that is displayed.

STEP 4

The SAR menu, showing the default options, will now appear on the console, and the cursor will be positioned to the first option. For each option, you must press ENTER to accept the default or overwrite the default with the desired value and press ENTER. If the response is valid, SAR will position the cursor to the next response. SAR will display a line of valid responses for the current option. The process option messages are detailed in [Section 15.04](#).

NOTE: SAR does not operate in true full screen mode. Do not use the cursor keys to move around the screen. Also the ERASE EOF key is not supported; if the remainder of a line is to be blanked out, it must be done using the space bar.

IT IS ALWAYS NECESSARY TO PRESS ENTER ONCE FOR EACH OPTION THAT THE CURSOR POSITIONS TO, WHETHER OR NOT THE DEFAULTS ARE ACCEPTED.

You can configure SAR to bypass some or all of the options, accepting the preset defaults without prompting. See member \$\$SAR in the FDR ICL (Installation Control Library) for details.

STEP 5

All devices needed for this execution should be available and ready. You may pre-mount Input or output tape volumes if you like, or SAR will prompt you to mount them when required.

STEP 6

The DUMP, RESTORE or CLIP will begin when all of the processing options have been specified. Successful conclusion of the requested function will be indicated by an 'FDR999 SAR SUCCESSFULLY COMPLETED' message. Other conditions are indicated by appropriate messages to the operator console and hardcopy device.

CONTINUED . . .

15.03 CONTINUED

**PROBLEM
DETERMIN-
ATION**

- If SAR does not successfully load, restart with Step 1.
- If SAR appears to load but does not display the SAR options on your console when you press ENTER, check the last 2 bytes of the PSW (you may need to STOP the CPU to see the PSW). If it ends in FFFF, either SAR is not receiving the ATTENTION from your console or it cannot identify it as a supported console type. If it ends in 3333, an unrecoverable error occurred on the console. If it ends in anything else, SAR identified another terminal as a console before you hit ENTER (possibly an end-user terminal).
- A runaway backup tape may be caused by an incorrect response to the INPUT TAPE UNIT= message.

**RESTART-
ING SAR**

Once a SAR function completes, either normally or abnormally, you may restart SAR to initiate a new function without re-IPLing. You will receive message:

FDR938 PRESS ENTER ON CONSOLE TO RESTART

In most cases, simply press ENTER on the SAR console to restart SAR; if that is not possible, use the PSW RESTART function (described below).

SAR will completely reinitialize itself and will redisplay the initial SAR menu. It will not retain any options you specified for the previous function (except for the console address); start over from step 4 to initiate the new SAR function.

CONTINUED . . .

15.04 SAR CONSOLE MESSAGES

**HARDCOPY
OPTIONS**

```

FDR - INNOVATION DATA PROCESSING - SAR VER 5.3/01P PRESS ENTER AFTER EACH REPLY
HRDCOPY DEVICE=1403, OVERRIDE=Y
HRDCOPY UNIT=00E
OPERATION REQUEST=RESTORE
TYPE=FULL
INPUT TAPE DEVICE=3480,1
INPUT TAPE UNIT=180,001
MODE=D4
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380
VOLUME SERIAL=XXXXXX, CPY=Y
FDR931 VALID RESPONSE=3211 1403 3203 4245 3286 3287 3288 3289 NONE TAPE 3480

```

The SAR console will look similar to the above after the IPL is complete or after SAR is restarted. The options shown are the defaults distributed by Innovation, but your installation may have changed some of the defaults with FDRSARLR (See Section 15.20) so your initial menu may look different. The options discussed on the following pages are highlighted in the menu at the top of each page. As shown, SAR will display a list of valid responses for most of the options. Don't forget that you must press ENTER after each line and let SAR reposition the cursor to the next option; do not attempt to move the cursor yourself.

HRDCOPY DEVICE= mmmm|**NONE, OVERRIDE=Y|C|N**

Indicate the type of hardcopy console or printer for SAR messages and diagnostic memory dumps.

DEVICE TYPE	HRDCOPY DEVICE=
Line printer	1403
3270-family printer	3287
Tape (cartridge, any type)	3480
Tape (9 track round)	TAPE
No hardcopy assigned	NONE

As you can see in the "valid responses" shown in the sample screen above, other device types are accepted, and you can enter your actual device if it is shown. However, the responses shown in this table cover all types of devices supported by SAR.

"Line printer" includes most high-speed printers that can be controlled by JES, except for 3800, 3900, and APF or page printers. "3270-family printer" includes printers that respond to 3270-family printer commands such as MVS hardcopy console devices.

If a hardcopy device is not available, reply 'NONE' to this message; you will not be prompted for the HARDCOPY UNIT. However, Innovation recommends that a hardcopy be assigned whenever a device is available, so that message listings and memory dumps can be preserved for diagnostic purposes.

If a tape is specified for the hardcopy device, SAR will write all messages to the tape drive specified. SAR will display the data set name and expiration date of the tape mounted and you will be asked if it is OK to write on the tape and given an opportunity to specify a new data set name. Care should be taken not to specify the FDR backup tape as hardcopy. SAR writes unblocked 120 byte records to the tape. IEBGENER, or any other print program, can be used to print the tape.

OVERRIDE= specifies whether the operator will be permitted to override the options that were set by FDRSARLR. Used primarily under VM. See member \$\$\$SAR in the FDR ICL (Installation Control Library) for details.

HRDCOPY UNIT=uuu

Provide the 3- or 4-digit hexadecimal I/O address of the hardcopy device. If NONE was specified for HRDCOPY DEVICE= this option will be bypassed (no reply required).

15.04 CONTINUED

OPERATION
OPTIONS

```

FDR - INNOVATION DATA PROCESSING - SAR VER 5.3/01P PRESS ENTER AFTER EACH REPLY
HRDCOPY DEVICE=1403, OVERRIDE=Y
HRDCOPY UNIT=00E
OPERATION REQUEST=RESTORE
TYPE=FULL
INPUT TAPE DEVICE=3480,1
INPUT TAPE UNIT=180,001
MODE=D4
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380

```

OPERATION REQUEST=RESTORE|DUMP|CLIP

RESTORE SAR is to restore a backup tape to a disk volume.

DUMP SAR is to backup a disk volume or data sets to tape.

CLIP SAR is to change the volume serial of a disk volume. When you press ENTER, the cursor will be positioned to the OUTPUT DISK UNIT response since the intervening options are not applicable. You will be prompted for the new volume serial after the disk's current volume serial is entered and validated. A tape is not used and the disk volume is not changed in any other way.

TYPE=FULL|DATASET|ABSTRK

FULL For a RESTORE, SAR is to restore all of the tracks contained on the backup tape.
For DUMP, SAR is to dump to tape all of the allocated tracks on the disk volume, including the VTOC and label track (cylinder 0 head 0).

DATASET Not valid for RESTORE. For DUMP, SAR is to prompt for the data set names or groups to be dumped. A maximum of 10 data set names or groups may be specified. A group name is specified by ending the characters with an *, requesting SAR to dump all data sets which begin with the characters specified. For example, DATASET SET NAME=SYS1.* will dump all data sets beginning with 'SYS1.'. The VTOC and cylinder zero head zero may be dumped by specifying ***VTOC. When all of the data set names have been entered, enter ***END to process.

ABSTRK Requests that specific tracks be dumped or restored. SAR will prompt you for the beginning address of a range of tracks with the message STARTING CCHH=CCCCCHHHH, and will position the cursor after the equal sign. Overtyping the CCCCCCHHHH with the five digit decimal cylinder number and four digit head (track) number of the first track in the range, and hitting ENTER. Next SAR displays the message ENDING CCHH=CCCCCHHHH; enter the decimal ending address of the last track of the range, and hit ENTER. SAR then returns to the STARTING CCHH message, and you enter the next range of tracks (maximum of 10). When all of the track ranges have been entered, enter ***END to process.

WARNING: Care should be taken if the volume label track or a track in the VTOC is to be restored. If cylinder zero, track zero is specified on a restore, the label will be restored from the dumped volume.

CONTINUED . . .

15.04 CONTINUED

TAPE
OPTIONS

```

FDR - INNOVATION DATA PROCESSING - SAR VER 5.3/01P PRESS ENTER AFTER EACH REPLY
HRDCOPY DEVICE=1403, OVERRIDE=Y
HRDCOPY UNIT=00E
OPERATION REQUEST=RESTORE
TYPE=FULL
INPUT TAPE DEVICE=3480,1
INPUT TAPE UNIT=180,001
MODE=D4
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380
VOLUME SERIAL=XXXXXX, CPY=Y
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380

```

If you have chosen OPERATION TYPE=DUMP, **INPUT TAPE** will change to **OUTPUT TAPE** on the SAR menu.

INPUT TAPE DEVICE=tttt,n

Specifies the type of device used for the backup tape.

DEVICE TYPE	SAR RESPONSE (tttt)
3420 Type tape drive (9 track round tape)	3420
3422 Type tape drive (9 track round tape)	3422
3480/3490 cartridge tape drive (18 track)**	3480 or 348X*
3490E cartridge tape drive (36 track)**	3490 or 349X*
3590 (Magstar) cartridge drive**	3590 or 359X*

* For a DUMP operation, SAR will invoke the IDRC hardware compaction feature of the cartridge drive if you specified 348X, 349X, or 359X, creating a compressed backup. If 3480, 3490, or 3590 is specified, the backup will not be compressed. For a RESTORE operation, IDRC compacted tapes are automatically supported as long as the tape drive hardware supports it; it does not matter which form of the device type you use.

** For a non-IBM tape drive, such as the StorageTek Timberline, Redwood, or 9840, specify the IBM tape type emulated by the drive.

Number of tape units (,n) -- specifies the number of tape units, 1 or 2, for the dump or restore. If "2" is specified, SAR will use 2 consecutively addressed tape drives, depending on the tape drive address specified for TAPE UNIT in the next response. Normally the second tape drive will be the next sequentially higher unit address (e.g., if you specify TAPE UNIT=391 the other unit will be 392). However, if the TAPE UNIT address ends with F, then the next lower address will be used (e.g., for TAPE UNIT=38F, the other unit will be 38E). This parameter reduces elapsed time because it allows the operators to pre-mount the next tape without waiting for the rewind and unload of the current tape. SAR will ask that the next input or output tape volume be mounted on the drive not in use. The mount requests will flip-flop until the dump or restore is completed. SAR will process the first tape on the unit specified in the TAPE UNIT parameter.

"n" can also be "R", e.g., "INPUT TAPE DEVICE=3590,R". This tells SAR to keep (retain) the tape at the end of a restore instead of unloading it. This can be used when the tape contains multiple backups to be restored with SAR, so that it doesn't have to be mounted multiple times.

CONTINUED . . .

15.04 CONTINUED

TAPE OPTIONS INPUT TAPE UNIT=uuu ,fff
(continued)

"uuu" is the 3- or 4-digit hexadecimal I/O address of the first or only tape drive (if the address is 4 digits, overtype the comma following "uuu" but be sure to keep the file number in the same location). If "2" was specified for the number of units in the reply for TAPE DEVICE, then a second tape drive with an address one higher or one lower than this drive will also be used, as described above.

"fff" is the file number on the input tape. It is used only for restore, and only when the tape is a multi-file tape (such as an ABR tape). The default is "001" which is used to restore from the first file on the tape. If you specify any other value (all 3 digits must be given), SAR will scan the labels on the tape until it finds the specified file number. For ABR tapes, the proper file number can be obtained from the FDR305 message produced when the backup was taken, or from the FDRABR "PRINT CATLG" report.

WARNING: We strongly recommend that ABR users run regular PRINT CATLG reports because if your system is down you may have no other way of identifying the required input tapes.

MODE=mm

Mode is meaningful primarily for tapes drives that can write in multiple densities, such as 3420 (round) tapes, and was used only for dumps since the tapes adjusted automatically to the density of tapes being read for restores. SAR will set "mm" to D4 for all cartridge drives, and to D3 (6250 BPI or the highest density the drive supports) for all round tape drives. You should not change this value.

DISK
OPTIONS

```
FDR - INNOVATION DATA PROCESSING - SAR VER 5.3/01P PRESS ENTER AFTER EACH REPLY
HRDCOPY DEVICE=1403, OVERRIDE=Y
HRDCOPY UNIT=00E
OPERATION REQUEST=RESTORE
TYPE=FULL
INPUT TAPE DEVICE=3480,1
INPUT TAPE UNIT=180,001
MODE=D4
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380
VOLUME SERIAL=XXXXXX, CPY=Y
OUTPUT DISK UNIT=130
OUTPUT DISK DEVICE=3380
```

If you have chosen OPERATION TYPE=DUMP, **OUTPUT DISK** will change to **INPUT DISK** on the SAR menu.

OUTPUT DISK UNIT=uuu

The 3-or 4-digit hexadecimal I/O address of the disk device.

CONTINUED . . .

15.04 CONTINUED

TAPE OPTIONS OUTPUT DISK DEVICE=dddddd**(continued)**

Specifies the type of disk SAR is to DUMP or RESTORE. For most disks, SAR will automatically determine and display the actual type of the disk after you reply to the DISK UNIT= option.

DEVICE TYPE	SAR RESPONSE
3380 Single Density (885 cylinders)	3380
3380-E Double Density (1770 cylinders)	3380-E
3380-K Triple Density (2655 cylinders)	3380-K
3390-1 Single Density (1113 cylinders)	3390-1
3390-2 Double Density (2226 cylinders)	3390-2
3390-3 Triple Density (3339 cylinders)	3390-3
3390-1/2/3 in 3380 Compatibility Mode	3380-1 or 3380-2 or 3380-3
3390-9 (10017 cylinders)	3390-9
IBM RAMAC, IBM RVA, EMC Symmetrix, StorageTek ICEBERG/ SVA and others	enter the name of device type and size emulated

NOTE: See "Disk Support" in [Section 15.01](#) for considerations when restoring to a device of the same type with a different capacity, such as 3390-2 to 3390-3.

VOLUME SERIAL=vvvvvv

"vvvvvv" is the current volume serial number of the disk being restored, dumped or clipped. SAR will verify that the correct volume is mounted; if you specified the wrong volser, SAR will replace it with the actual volser. **For RESTORE, if the volume does not contain a valid label or if the volume is brand new, type in six blanks or an asterisk (*) and SAR will bypass the label check. For CLIP, SAR will prompt you to enter the new volser.**

CPY=Y|N|C

Specifies the processing of the disk volume serial during RESTORE.

Y - the restored volume will receive the volume serial of the dumped volume, from the backup tape.

N - the volume serial of the receiving volume is to be retained unchanged.

C - the volume serial is to be changed to a user-specified value. SAR will prompt you for the new volume serial.

**ENDING
MESSAGES**

The SAR operation will begin after all of the above messages are satisfied. On a DUMP operation, SAR will display additional messages for the data set name and use of the scratch tapes. Successful completion is indicated by the FDR999 message on the hardcopy device and the operator console. The number of tracks processed is also displayed. SAR can be restarted to perform another function by pressing ENTER on the console.

CONTINUED . . .

15.20 STAND-ALONE LOADER UTILITY

The FDR Stand-Alone loader utility (FDRSARLR) can be used to:

- construct a new IPLable copy of SAR. This SAR can be written to the label track of a disk volume, to unlabeled or standard labeled tape, or can be output in card-image format to any sequential data set.
- modify an IPLable copy of SAR on the label track of a disk volume.
- modify the defaults for SAR processing (menu) options in the new or updated SAR.
- apply INNOVATION-supplied fixes ("zaps") to the new or updated SAR.
-

This utility will print a processing summary map upon completion. This map will indicate the location of the SAR program, console options, message defaults, and loader utility options in effect.

**EXECUTING
SAR FROM
TAPE**

The SAR program may be IPLed from unlabeled (NL) tape or standard labeled (SL) tape. The FDR distribution tape is an SL tape with the first file containing a loadable SAR program.

When IPLing from SL tape, the initial program load function (IPL) must be performed five times to bypass the labels. An NL tape requires only one IPL, so it can be loaded much faster. You may create an unlabeled tape containing SAR by using the Stand-Alone loader utility (FDRSARLR), or by copying the first file of the distribution tape (DSN=SAR) to an unlabeled tape with IEBGENER, using DCB attributes of (RECFM=F,BLKSIZE=80,LRECL=80). FDRSARLR must be used if you wish to change options or apply fixes.

**EXECUTING
SAR FROM
DISK**

SAR may be IPLed from disk, using FDRSARLR to write an IPLable copy of SAR onto cylinder 0, track 0 (the label track) so it occupies no space on the volume. SAR can be written to any disk that does not already contain IPL text (such as the MVS SYSRES volume or a volume containing the IBM stand-alone memory dump (SAD)). The FDRSARLR will not let you accidentally overwrite other IPL text (Contact Innovation for instructions if you WANT to overwrite other IPL text on a disk volume with the SAR program).

**CONTROL
STATEMENTS**

SAR is distributed with defaults for all options except console addresses. If these defaults do not match your environment all console options and option defaults can be supplied via loader utility control statements. The loader will change the output copy of SAR to have the new defaults.

All option defaults specified on FDRSARLR control statements may be overridden at IPL time,. However, if you wish to accept certain defaults with no possibility of override, see member \$\$\$SAR in the FDR ICL (Installation Control Library) for details. This is especially useful under VM where you may wish to have multiple SAR decks with all options preset, so that SAR starts executing immediately without operator input.

FDRSARLR control statements may also be used to apply fixes to SAR.

**STAND-ALONE
LOADER
SECURITY
OPTIONS**

FDRSARLR supports Security checking if OBJIN or OUTPUT points to a disk volume (see Section 90).

If the ALLCALL option has been enabled, a SAF call for class DASDVOL will be issued, ACCESS=READ for OBJIN, ACCESS=ALTER for OUTPUT.

15.21 FDRSARLR JOB CONTROL REQUIREMENTS

**EXEC
STATEMENT**

Must specify PGM=FDRSARLR and may need to specify REGION=256K or more. You normally don't need to specify a PARM= operand, but you may specify one or more of the following PARM= options. If more than one operand is specified, the operands must be enclosed in parentheses or apostrophes and separated by commas, e.g., PARM='REWRITE,PRSYSIN'. Operands may be specified in any order.

REWRITE - Only meaningful if the OUTPUT DD statement specifies a disk. Directs FDRSARLR to write the SAR program to the label track of a volume which already contains a IPLable version of SAR. REWRITE is assumed if both OBJIN and OUTPUT point to the same disk (which allows you to read SAR from a disk, update it with options or fixes, and write it back).

ERASE - Only meaningful if the OUTPUT DD statement specifies a disk. Directs the loader utility to remove a copy of SAR from the label track of the volume. The volume will no longer be IPLable. OBJIN is not required.

If neither REWRITE nor ERASE is specified, and the OUTPUT DD specifies a disk, FDRSARLR will write an IPLable copy of SAR to the label track of that disk only if it currently does not contain any IPLable program.

DECK - Only meaningful if the OUTPUT DD statement specifies a tape. Directs the loader utility to write the IPLable SAR program as fixed-length 80 byte records. DECK is assumed if the OUTPUT DD statement specifies a unit record device, a SYSOUT data set or a sequential disk data set. If DECK is not specified FDRSARLR will write SAR to tape in a tape-oriented format with just two records; this will usually speed up IPL from tape.

PRSYSIN - Directs FDRSARLR to print all records read from SYSIN. If omitted, only records in error are printed.

**STEPLIB OR
JOBLIB DD
STATEMENT**

If required, must specify the load module library in which FDR resides. This must be an authorized library.

**SYSPRINT DD
STATEMENT**

Specifies the primary output message data set. This is a required DD statement and is usually a SYSOUT data set.

**SYSUDUMP DD
STATEMENT**

Specifies the ABEND DUMP data set. Usually a SYSOUT data set.

**OBJIN DD
STATEMENT**

Specifies the input to the loader utility as either:

* A non-temporary data set containing a copy of SAR to be copied/modified. This may be file 1 of FDR distribution tape (DSN=SAR), member SAROBJ of the FDR Installation Control Library (ICL), or any other IPLable copy of SAR that you have created.

* A disk volume to which FDRSARLR has previously written an IPLable version of SAR (do not specify DSN=). In this case, the OBJIN and OUTPUT DD statements may specify the same disk volume, providing a convenient method to change the console options or processing option defaults.

OBJIN is not required if PARM=ERASE is specified.

15.21 CONTINUED

**OUTPUT DD
STATEMENT**

Specifies the output device for the loader utility as either a 9345, 3380 or 3390 disk volume (or any disk emulating one of these, such as IBM RAMAC), a tape drive containing a labeled or unlabeled tape, a card punch, a SYSOUT data set that will be directed to a card punch or a sequential disk data set.

If it points to a disk volume, and that disk already contains IPL text (such as the MVS sysres volume, or a volume containing the MVS stand-alone memory dump program (SAD)), FDRSARLR will not overwrite that IPL text. If the disk already contains a copy of SAR, it will not overwrite it unless you specified the PARM=REWRITE option, or OBJIN and OUTPUT point to the same disk volume.

**SYSIN DD
STATEMENT**

Optional. If present, specifies a data set containing control statements as described in the next section.

CONTINUED . . .

15.22 FDRSARLR CONTROL STATEMENTS

**CONTROL
STATEMENT
FORMAT**

The control statements may start in any column from 1 to 60. If keywords contain multiple words (Ex: CONSOLE DEVICE=), exactly one blank must be specified between words. A control statement may not be continued. Only one keyword may be specified per statement, even though multiple keywords may appear on the same line when the SAR runs. An asterisk (*) in column 1 indicates a comment.

If FDRSARLR determines that a control statement contains an error, the control statement and an error message will be printed. Other control statements will be checked for validity but no output will be written to DD OUTPUT.

**CONSOLE
CONTROL
STATEMENTS****CONSOLE DEVICE=dddd**

Specifies the device type of the console to be used by SAR. If not specified, SAR will determine at IPL time the type of console. You will probably be using display-type consoles; if so specify "3270". If you are using a printer-keyboard console (possibly under VM), specify "3215". If you are using the SCLP console (see [Section 15.02](#)) specify "SCLP".

CONSOLE UNIT=uuu

Specifies the 3- or 4-digit hexadecimal address of the console to be used by SAR. If not specified, at IPL time SAR will accept the first attention interrupt as indicating the console that it should use. You may specify up to five CONSOLE UNIT= statements. At IPL time, SAR will attempt to communicate with each address in the order that the CONSOLE UNIT statements appeared, and will use the first address that appears to be a console. If "CONSOLE DEVICE=SCLP" was specified, no CONSOLE UNIT= statements are required (the SCLP console does not have a device address).

**CONSOLE
MESSAGE
DEFAULTS**

You may enter any or all of the following statements, in any order, to change the defaults for SAR options in the copy of SAR being created. Changing these options to the values you will usually use will save time when initializing SAR. A complete explanation of these options and valid responses can be found in [Section 15.04](#), SAR Console Messages. FDRSARLR does **not** validate the response values; it simply stores them for SAR's use.

HRDCOPY DEVICE=mmmm**OVERRIDE=x****HRDCOPY UNIT=uuu****OPERATION REQUEST=mmmmmmm****TYPE=mmmm****INPUT TAPE DEVICE=tttt,n**

or

OUTPUT TAPE DEVICE=tttt,n**INPUT TAPE UNIT=uuu,fff**

or

OUTPUT TAPE UNIT=uuu,fff**MODE=mm****OUTPUT DISK UNIT=uuu**

or

INPUT DISK UNIT=uuu**OUTPUT DISK DEVICE=dddddd**

or

INPUT DISK DEVICE=dddddd**VOLUME SERIAL=vvvvvv****CPY=r**

15.23 FDRSARLR EXAMPLES

**LOAD SAR TO
DISK VOLUME**

Execute FDRSARLR to read SAR from the first file of the FDR distribution tape, specify the console address, change a console message default, and place SAR onto the label track of a disk volume. If you want to copy SAR to a disk without changing any options, omit SYSIN and the control statements.

```
//SATODISK EXEC PGM=FDRSARLR,REGION=256K
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//OBJIN DD UNIT=TAPE,VOL=SER=FDR53P, FDR53T if Trial
// DSN=SAR,DISP=OLD
//OUTPUT DD UNIT=3380,VOL=SER=vvvvvvv,DISP=SHR
//SYSIN DD *
        CONSOLE DEVICE=3270
        CONSOLE UNIT=005
        OPERATION REQUEST=DUMP
/*
```

NOTE: Specifying a console message control statement changes the default that appears on the console.

**LOAD SAR TO
DISK ALL
OPTIONS
SPECIFIED**

Execute FDRSARLR to read the copy of SAR in the FDR ICL (Installation Control Library) and place it onto the label track of a disk volume. All possible processing options are specified here as an example of their format. All control statements will be echoed on SYSPRINT.

```
//SATODISK EXEC PGM=FDRSARLR,REGION=256K,PARM='PRSYSIN'
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//OBJIN DD DISP=SHR,DSN=IDP.ICLFDR53(SAROBJ)
//OUTPUT DD UNIT=3390,VOL=SER=vvvvvvv,DISP=SHR
//SYSIN DD *
        CONSOLE DEVICE=3270
        CONSOLE UNIT=0E2
        CONSOLE UNIT=0F2
        HRDCOPY DEVICE=1403
        OVERRIDE=Y
        HRDCOPY UNIT=00F
        OPERATION REQUEST=RESTORE
        TYPE=FULL
        INPUT TAPE DEVICE=3480,2
        INPUT TAPE UNIT=390,001
        MODE=D4
        OUTPUT DISK UNIT=420
        OUTPUT DISK DEVICE=3390-3
        VOLUME SERIAL=MVSRES
        CPY=Y
/*
```

15.23 CONTINUED

**LOAD SAR TO
TAPE**

Execute FDRSARLR to place SAR onto an unlabeled tape volume. IPL from an NL tape is much more convenient than from a labeled tape.

```
//SATOTAPE    EXEC    PGM=FDRSARLR,REGION=256K
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//OBJIN       DD      UNIT=TAPE,VOL=SER=FDR53P,      FDR53T if Trial
//              DSN=SAR,DISP=OLD
//OUTPUT      DD      UNIT=TAPE,LABEL=(,NL),VOL=SER=IPLSAR,
//              DSN=IPLSAR,DISP=(,KEEP)
```

To place SAR onto a standard labeled tape volume, use the same JCL, except change the last two lines as shown below. This is not recommended, since you will have to perform the IPL function five times.

```
//OUTPUT      DD      UNIT=TAPE,VOL=SER=volser,
//              DSN=IPLSAR,DISP=(,KEEP)
```

**MODIFYING
SAR ON DISK**

Execute FDRSARLR to modify a copy of SAR previously written to disk and write it back to the same (or a different) disk. The example shows VER/REP statements used to apply an Innovation-provided fix to SAR, but this can also be used to change console addresses or message option defaults stored in that copy of SAR.

```
//SAZAP        EXEC    PGM=FDRSARLR,REGION=256K
//SYSPRINT     DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//OBJIN        DD      UNIT=3390,VOL=SER=xxxxxx,DISP=SHR
//OUTPUT       DD      UNIT=3390,VOL=SER=xxxxxx,DISP=SHR
//SYSIN        DD      *
                VER    aaaa dddd    VERIFY 'dddd' at address 'aaaa'
                REP    aaaa eeee    REPLACE 'eeee' at address 'aaaa'
                DUMP                                Print SAR object module
```

**ERASING SAR
FROM DISK**

Use FDRSARLR to erase a copy of SAR that was previously placed on a disk. ERASE verifies that the label track of that disk does contain SAR, so it will not erase any other IPL text.

```
//SAERASE      EXEC    PGM=FDRSARLR,REGION=256K,PARM=ERASE
//SYSPRINT     DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//OUTPUT       DD      UNIT=3390,VOL=SER=xxxxxx,DISP=SHR
```

CONTINUED . . .

15.24 FDRSARLR ISPF INTERFACE

The FDR ISPF install dialogs include a function for generating FDRSARLR jobstreams. Enter "A.I.2" to reach the FDRSARLR dialog (if "A" is not the character used in your installation to reach the FDR dialogs from the ISPF main menu, substitute the correct character).

PANEL A.I.2

```

----- SAR LOADER UTILITY -- ENVIRONMENTAL SPECIFICATIONS -----
COMMAND ==> SEL

      SEL   - SELECT SAR DEFAULTS PANEL      HELP   - TUTORIAL HELP PANELS

OUTPUT DISK DEVICE: (3380/3390)
  VOLUME SERIAL   ==>                               REWRITE   ==> NO       (YES/NO)

INSTALLATION CONTROL LIBRARY DATA SET:
  DATA SET NAME   ==> 'IDP.ICLFDR53'

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME   ==> 'IDP.MODFDR53'

SYSOUT CLASS      ==> *

JOB STATEMENT INFORMATION:
  ==> //userA JOB (ACCOUNT),'NAME',
  ==> //          NOTIFY=user
  ==> //*                               OUTPUT DISK DEVICE=3380

```

On this panel, enter the volume serial of the disk volume to which SAR is to be written. You may also need to update the names of the FDR ICL and program libraries, and specify proper JOB statement information. When done, enter SEL on the command line to proceed to the next panel.

```

----- SAR LOADER UTILITY -- SELECT SAR DEFAULTS -----
COMMAND ==>

      EDIT   - EDIT SAR LOADER JCL           SUBMIT - SUBMIT SAR LOADER JOB
      HELP   - HELP TUTORIAL PANELS         CANCEL - EXIT IMMEDIATELY

CONSOLE DEVICE.... ==>
CONSOLE UNIT..... ==>                               (UP TO 5 UNITS)
OPTIONS THAT CAN BE OVERRIDDEN AT IPL TIME:
HRDCOPY DEVICE.... ==> 1403
HRDCOPY UNIT..... ==> 00E
OPERATION REQUEST. ==> RESTORE
TYPE..... ==> FULL
TAPE DEVICE..... ==> 3480
TAPE UNIT..... ==> 180,001
MODE..... ==> D4
DISK DEVICE..... ==> 3380
DISK UNIT..... ==> 130
DISK VOLUME SERIAL =

```

Here you may specify all of the options for customizing SAR. Make any desired changes to the default values shown (any changes you make will be remembered and will be redisplayed if you use this dialog again in the future) or leave them alone to accept the defaults. All the option values shown will be passed to FDRSARLR (even those that you did not change).

On the command line, enter SUBMIT to directly submit the batch job to copy SAR to the designated disk volume, or EDIT to edit the jobstream. From EDIT you can modify the jobstream (e.g., to output to tape) or save it for future execution.

15.25 SAR HINTS AND RECOMMENDATIONS



► IPL from Disk

We strongly recommend that copies of SAR be put onto disk. The SAR IPL records can be loaded onto disk by the FDRSARLR program. SAR will reside on the volume label track, cylinder 0 track 0, and will not take any space that could be allocated to data sets. We recommend putting copies of SAR onto several volumes on different control units. Critical time can be lost when SAR is needed and the tape cannot be found, or the only copy is locked in the system programmer's desk. The FDR Installation Control Library (ICL) contains a member called SAROBJ which may be used as input to FDRSARLR to create a copy of SAR on either tape or disk. Of course, keep a visible list of the volumes on which SAR has been loaded near the console.

While tape will always work, you have to find the tape first!!

► Customizing SAR

You can use the FDRSARLR program to customize SAR for your installation (or for your disaster recovery site). The options that are displayed on the SAR menu are clear and can be overridden at run time, so they usually do not **need** to be customized, but it may save time to set some options to the commonly used values.

However, it is often very useful to customize SAR with the console device type and console address (or addresses, up to 5). Why? If the console address is not specified in advance by FDRSARLR, SAR waits for a unit to present an ATTENTION interrupt (which is generated by a console when the ENTER or REQUEST key is pressed). SAR then issues CCWs for various type of consoles to see if this device responds to one of them. However, if you have terminals which are attached to your CPU via local non-SNA control units, anxious users pressing ENTER may receive the SAR menu instead of you. If you customize SAR with the addresses of the consoles you normally use, SAR will test them immediately without waiting for an ATTENTION.

If you are reading this, your system is down, and you do not get the SAR menu after IPLing SAR, you probably have terminals or other devices which are presenting ATTENTION. You may have to reset and/or disable the control units of those device via switches on their control panels before retrying SAR. You may also need to disable or reset other communication controllers (such as 3745s) and CTCs (channel-to-channel adaptors) as these may also generate ATTENTION.

► Eliminate tape rewind delays

SAR supports using two tape units for dump and restore. Using two tape drives will reduce elapsed time by making it possible for the operator to premount the next input or output tape while the current tape is being read or written. If the next tape is mounted and readied on the inactive drive, SAR will automatically start using it when the current tape volume is done (a DUMP will require a confirmation on the console before writing to the tape). However, you must be careful to mount the tapes on the right drives in the right order.

When restoring more than one disk volume with SAR, typical at a disaster recovery site, the elapsed time savings will be significant, about 25% reduction average.

You can specify the two tape options at run time, or customize SAR to use it by default with FDRSARLR, changing:

```
INPUT TAPE DEVICE=3480,1
to
INPUT TAPE DEVICE=3480,2
```

15.25 CONTINUED

► Retain tape between restores

Modern high-speed, high-capacity tape drives such as the IBM Magstar 3590 and the StorageTek Redwood and 9840 can hold the backups of many disk volumes. Most installations will stack multiple backups on tape to get the best utilization from their tape cartridges; ABR backups are automatically stacked on tape.

Normally SAR will rewind and unload each input tape when it completes the restore.

If you are restoring from a backup tape containing multiple backups, and plan to restore another backup on the same tape, you can specify that SAR should NOT unload the tape at the end of the restore, by:

```
INPUT TAPE DEVICE=3580,R
```

► ABR backup tapes

SAR supports restores from any file on an ABR full-volume backup. To be prepared to run SAR to restore ABR backups, you should run a weekly FDRABRP "PRINT CATLG" report, which will show the tapes volumes and file numbers for the full-volume backups of every disk volume. These hardcopy reports should be readily available to the Operations staff so that the proper tapes can be fetched and mounted. If you need to run SAR and the PRINT CATLG report is not available, you can also get the same information from the FDR305 messages in the listing of the last FDRABR job that created the full-volume backups required (but only if the printout is available in hardcopy!).

You must specify the file number as "fff" (specify all 3 digits) in the SAR option: INPUT TAPE UNIT=uuu,fff

The number that you specify is the FILE value from the PRINT CATLG or FDR305 message, as shown below. This is the logical file sequence number (relative to the first volume in a multi-file, multi-volume tape aggregate); it may not be the physical file number on the current tape. SAR will match the "fff" value against the header labels of the tape mounted until that file is found.

SAMPLE PRINT CATLG REPORT

FDRABR VOLUME BACKUP REPORT															
VOLSER	GEN	CYCLE	TYPE	DUMP	DATE	TAPE	FILE	DATA	SET	NAME	COPY	FILE	TAPE	VOLUME(S)	REQUIRED
IDPPX0	143	00	FDR		94.294	FDRABR.VIDPPX0.C1014800					1	5	B80667,B80670		
		00	FDR		94.294	FDRABR.VIDPPX0.C2014800					2	5	BV1060,BV1066		
		01		DSF	94.297	FDRABR.VIDPPX0.C1014801					1	9	B80689		
		02		DSF	94.298	FDRABR.VIDPPX0.C1014802					1	19	B80693		

SAMPLE ABR MESSAGES

```
FDR304 FDR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=IDPPX0,UNIT=3390-2
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VIDPPX0.C1014800,FILE=005 VOL=SER=B80667 B80670
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VIDPPX0.C2014800,FILE=005 VOL=SER=BV1060 BV1066
FDR306 DUMP SUCCESSFULLY COMPLETED
```

CONTINUED . . .

15.25 CONTINUED

► VOLUME SERIAL reply

Sometimes users are not sure what volume serial to specify for a SAR restore. The correct serial is the current volume serial of the output volume to which you are restoring. SAR will verify that volume serial before allowing the restore to proceed. If you enter the wrong serial, SAR will display the actual serial for verification.

If you specify the SAR option CPY=Y, then SAR will change the volume serial of that disk to the disk serial contained on the backup tape being restored (if different). If CPY=N is specified, then the serial of the output disk will be retained. If you specify CPY=C, SAR will return to the VOLUME SERIAL= prompt to ask for a new volume serial for the output disk. However, if the backup being restored contains VSAM datasets or is an SMS volume, the disk may not be usable by MVS if CPY=Y is not used (unlike FDR, SAR does not check for these conditions). CPY=Y is usually the appropriate choice.

For example, if you backed up volume SYS001 to tape and you are now restoring to volume serial SPARE1, you specify

VOLUME SERIAL=SPARE1,CPY=Y

and the volume will be labeled SYS001 after the restore is complete.

If the output volume does not contain a valid volume label or is brand new (uninitialized), enter 6 blanks or one asterisk (*) for the response to the volume serial prompt and SAR will allow the restore without verifying the serial. This allows you to do restores without running ICKDSF to preinitialize the volumes. If the volume contains no volume label, you may get a SAR I/O error message; it can be ignored by pressing ENTER to continue.

► Stand-Alone Backup

Unlike DFDSS, SAR can do stand-alone backups. Stand-alone backup allows more flexibility in a disaster situation. For example, if the operating system is unavailable due to an HDA failure, and you have no spare volumes to restore required system volumes to, you can use SAR to dump a less critical volume before overwriting it with the backup of the critical volume. You can restore the backup of the less critical volume later when the HDA problem has been repaired (be sure to carefully record the serials of the tape volumes used for the backup). Without stand-alone backup, someone may have to decide what volumes can be sacrificed in order to get the system back, or wait for the repair to be completed (which may be hours).

Even if spare volumes are available, can you be sure that no one has used it for "short-term" storage? To be sure, you can do a stand-alone backup of the spare volume before you do that critical stand-alone restore.

► Problem Determination

If the IPL of SAR fails, it may be a hardware problem or a configuration problem. To quickly determine if it is a hardware problem, try to IPL stand-alone ICKDSF from tape, or IBM's stand-alone memory dump (SADMP) from disk. However, both of these require preparation while your MVS system is still up.

If the IPL of SAR completes (IPL COMPLETE or LOAD COMPLETE messages) but the SAR menu does not appear on your console when you press ENTER (or on the console you configured as the default console for SAR), display the current PSW (you may have to STOP the CPU to see the PSW). If it does not end in FFFF, then there may be a hardware problem with the console, or it may be that another locally-attached terminal has pressed ENTER before you did and has received the SAR messages. If it does end in FFFF, then SAR is not recognizing your console as a supported SAR console.

CONTINUED . . .

15.25 CONTINUED

► IPL From Standard Label Tape

There is often confusion when you IPL SAR from a SL tape, such as the FDR distribution tape. This is why we recommend that you copy SAR to a NL (unlabeled) tape or disk volume to make IPL simpler. But if you must IPL from an SL tape:

The first 4 IPL attempts will fail. This is because they encounter the VOL1, HDR1, HDR2 tape labels, and the tape mark that follows them. Each IPL moves past one record or tape mark. The 5th IPL will read the SAR IPL text in the data file and will be successful.

Depending on your CPU type, you will get various error messages when those first 4 IPLs fail. If your CPU displays the CSW or status, the channel status will probably be 0E00 or 0200 for the first 3 attempts (unit check) and 0D00 or 0100 for the 4th (unit exception, caused by the tape mark). On some CPUs you must allow 10-15 seconds between IPL attempts; if you attempt to relIPL too quickly, some CPUs will fail the IPL attempt without moving the tape.

SAMPLE SAR SCREEN AFTER RESTORE IS COMPLETE

```
FDR - INNOVATION DATA PROCESSING - SAR VER5.3/01P PRESS ENTER AFTER EACH REPLY
HRDCOPY DEVICE=1403, OVERRIDE=Y
HRDCOPY UNIT=00E
OPERATION REQUEST=RESTORE
TYPE=FULL
INPUT TAPE DEVICE=3480,2
INPUT TAPE UNIT=5A0,005
MODE=D4
OUTPUT DISK UNIT=CF2
OUTPUT DISK DEVICE=3390-2
VOLUME SERIAL=IDPPX0, CPY=Y
-FDR933W TAPE MOUNTED FOR RESTORE IS VOL=B80667 DSN=.VIDPPX0.C1014800
-FDR933W TAPE MOUNTED FOR RESTORE IS VOL=B80670 DSN=.VIDPPX0.C1014800
-FDR999W SAR RESTORE SUCCESSFULLY COMPLETED VOL=IDPPX0 NVOL=IDPPX0 TRACKS=007485
FDR938 PRESS ENTER ON CONSOLE TO RESTART
```

20.01 DSF TECHNICAL SUMMARY

FDRDSF Program FDRDSF (Data Set Functions) performs backups and restores of selected data sets and VSAM clusters. It can also do data set restores from full-volume backup tapes produced by FDR, ABR or SAR. [See Section 21](#) for a description of FDRCOPY, which performs disk-to-disk data set copies.

DATA SET BACKUPS FDRDSF backups are similar to FDR backups, using the same format, but containing only the information and data tracks associated with the data sets selected by DSF control statements. You may select data sets by specific name, by prefix, or by sophisticated data set name masking. DSF backups operate on one DASD volume at a time, reading the VTOC of the volume, backing up the data sets you have selected and writing the backup to a tape or DASD data set; a separate backup data set is required for every DASD volume processed but that backup data set may contain backups of many original DASD data sets. The backup will contain an image of:

- all tracks which are allocated to the selected data sets and VSAM clusters, according to the DSCBs in the VTOC.
- an edited version of information from the VTOC and VVDS is also stored at the beginning of the backup, serving as an index of the data actually in the backup, and providing information for allocation of output data sets.
- The VTOCIX and VVDS may be backed up as data sets, if DSF control statements select them. However, none of these are required to be able to restore data sets from the backup, and they are usually not backed up.
- some information contained only the catalog entry for the data sets being backed up will also be included in the backup. This currently includes PATH definitions for the alternate indexes (AIXs) of ICF VSAM clusters and the aliases of ICF VSAM user catalogs. Other catalog-only information, such as OWNERID, is not backed up and so will not be recreated during a data set restore.

20.01 CONTINUED

**DATA SET
RESTORES**

FDRDSF data set restores are considerably different from FDR full-volume restores:

- you may restore all data sets on the backup, or only selected data sets and/or VSAM clusters. The input to the restore may be a FDR, ABR or SAR full-volume backup, or a DSF or ABR data set backup.
- the data tracks of the selected data sets may be restored to a different physical location (cylinder and head address) than the original tracks occupied. In other words, you do not have to be concerned about the location of the output data set or the number of extents it is in.
- data sets and clusters may be renamed during restore.
- if an output data set (either the original data set name or the new name) already exists, and is large enough, DSF will replace the contents of the existing data set and update its describing information (VTOC and VVDS information). By default, DSF will restore an existing data set to whatever volume it is currently cataloged.
- if an output data set does not exist, DSF will allocate it, automatically making it large enough to hold the contents of the input data set. Usually DSF will also catalog the output data set as well (if the output data set is already cataloged to another volume, DSF will not catalog it unless you instruct it to do so). VSAM clusters must be cataloged when they are allocated, so a VSAM allocation will fail if it is already cataloged (but you have an option to delete the existing cluster and catalog the new one in its place).
- DSF usually restores data sets to the same device type they were backed up from (e.g., from a 3390 backup to a 3390). This is called a like device or physical restore. The size (number of cylinders) on the original disk volume and the output volume is not important, as long as the output has enough free space to hold the output data sets being allocated. In a like restore, the original data tracks of the selected data sets are restored exactly as they were backed up (but there is an option to reblock certain data set types).
- DSF can also restore data sets to a different device type (such as a backup of a 3380 to an output 3390), with some restrictions. This is called an unlike device or logical restore since the data records of the original data sets are usually reformatted to make better use of the track capacity of the output disk.
- Output data sets may be directed to various output disk volumes. Data sets from one backup, originally all on the same disk, can be restored to many output disks concurrently (reading the backup only once). FDRDSF JCL must point to the backup data sets, but the output volumes can be identified by JCL or by DSF control statements. You can also identify a list or group of volumes as the target volume; DSF will find one with sufficient space for the data set.
- Multi-volume data sets can be restored, but only to the same number of volumes they originally occupied when dumped. Multi-volume VSAM is handled, but only when restored to the original device type.
- At the end of the restore, DSF will update the DSCB of the output data set and, for VSAM and SMS-managed data sets, its VVDS entry, so that they properly describe the data that was restored. If a restore is interrupted or cancelled, this update will not be done and the data sets will probably be unusable even though all data tracks were restored.
- If DSF allocated an output data set and the restore of that data set gets errors, such as disk I/O errors, the data set will be deleted from disk, to avoid leaving unusable and uncataloged data sets.

20.01 CONTINUED

**CATALOGING
NON-VSAM
OUTPUT DATA
SETS**

When restoring single-volume non-VSAM data sets:

If DSF allocates space for the output data set, then the default is that DSF will catalog the data set to the output volume, unless the data set was already cataloged. So, if the data set already exists on another volume and is cataloged to it, a DSF restore will not disturb it.

If the RECAT operand is specified, it will catalog the newly allocated data set to the output volume, whether the data set was cataloged to the input volume before the operation, or cataloged elsewhere, or not cataloged at all. This allows you to insure that the catalog will point to the restored data set.

If the NOCAT operand is specified then DSF will not catalog the newly allocated data set. This is ignored for SMS-managed data sets which must always be cataloged.

Cataloging or recataloging of non-VSAM data sets occurs at the end of the restore, and is bypassed if any restore errors have occurred for a given data set.

If the output data set exists on disk before the restore, then DSF will not update the catalog, unless the CATIFALLOC operand is specified. CATIFALLOC causes DSF to apply the same rules described above for data sets allocated by DSF.

STPCAT DDs are supported; if present, only the STPCAT catalog (or the first catalog in the STPCAT concatenation) will be searched and updated for non-VSAM data sets. The target catalog for VSAM clusters is controlled by the ICFCAT= operand described in [Section 20.08](#). STPCAT is not supported if SMS-managed data sets must be allocated.

NOTE: If a DSF restore is interrupted by an ABEND or system crash, then the output data sets will be left on disk, but will not be cataloged. You should resubmit the restore, to insure that all data tracks are restored, but the output data sets will already be allocated so DSF will not catalog them unless you specify the operand CATIFALLOC.

For multi-volume non-VSAM data sets, the above rules for single-volume data sets apply, with the following modifications:

If the data set is not already cataloged DSF will create a new multi-volume catalog entry with the current output volume in the proper slot as indicated by the volume sequence number in the DSCB. If the volume sequence number is higher than 1, FDRCOPY will fill in the preceding slots in the catalog entry with a dummy volume serial of "#####nn".

If the data set is already cataloged DSF will update the catalog entry by putting the current output volume into the proper slot as indicated by the volume sequence number in the DSCB. On a RESTORE without RECAT, DSF will update the catalog entry only if the slot for this output volume contains the dummy volume serial of "#####nn"; otherwise a warning message will be issued.

So, when all pieces of a multi-volume data set have been restored, the catalog entry will properly point to all the volumes to which it was restored (remember that a multi-volume data set must be restored to the same number of volumes it was backed up from). If the data set is already cataloged, this is true only if RECAT is specified. If you do not restore one or more pieces of the data set, the catalog entry will be inaccurate and the data set will not be usable (it may contain #####nn or original volsers). It is your responsibility to insure that the data set is backed up from all the volumes it occupies and is restored to the same number of volumes.

All of these rules sound complicated, but they are designed to automatically do the right things when cataloging non-VSAM data sets. You need to be concerned only if a data set being restored may already be cataloged to another volume; if so, specify RECAT if you want the newly restored data set to be the cataloged data set, leaving the other data set uncataloged. Specify CATIFALLOC if there is a chance that the output data set is already allocated but is not currently cataloged (and you want it to be the cataloged version).

20.01 CONTINUED

**ICF VSAM
SUPPORT**

DSF supports backup, restore, and allocation of ICF VSAM clusters similarly to the way it handles non-VSAM data sets, with these differences:

- DSF always refers to all components of a cluster by the base cluster name. You cannot select clusters by component name.
- For clusters that have alternate indexes (AIXs), the alternate indexes are automatically selected, using the base cluster name. You never refer to an AIX by its AIX name.
- DSF backup will select all components of the selected clusters (including AIX components) from the volumes being backed up. If a cluster has components on multiple volumes, it is your responsibility to backup that cluster from all of its volumes (you may want to use ABR Application Backup, described in [Section 52](#) to automate this type of backup).
- If DSF restore finds that the cluster's components already exists on the output disks, it will restore the data back to the existing allocation.
- If the cluster doesn't exist, DSF restore will allocate and catalog the cluster at the beginning of the restore (unlike non-VSAM data sets, VSAM clusters must be cataloged at the time they are allocated). If the backup contains a base cluster and one or more of its AIXs, the base will be defined before the AIXs automatically. But if the base cluster is on one volume and the AIX on another, it is your responsibility to restore the base cluster before the AIXs; contact Innovation for assistance if necessary.
- PATHs are catalog-only entries which relate base clusters to AIXs. When DSF defines an AIX, it will also define its PATHs, but only if both the backup and restore were done with V5.3 level 30 or above.
- If the cluster does not exist on the output volume but is already cataloged, the allocation will fail unless you specify the VRECAT operand. VRECAT allows DSF to delete the cataloged cluster and redefine the new one in its place; this works even if only the catalog entry for the cluster exists (no actual data set on another disk).
- If the cluster includes a multi-volume component, or has data and index on separate volumes, special procedures are used ([See Section 80.13](#)). The cluster will not be usable until all components have been restored.
- On a like device restore, data tracks in the cluster are restored exactly as they were when dumped; the cluster is not reorganized. Single-volume clusters can be restored to an unlike device (e.g., 3380 to 3390); although the data is rearranged to fit the new data tracks, it is not reorganized.
- at the end of the restore, the DSCB and VVDS entries for each component are updated.

[See Section 80.13](#) for VSAM Special Considerations and more information on VSAM operations.

**ABSOLUTE
TRACK
OPERATIONS**

FDRDSF can also dump and restore data tracks by their physical, absolute track address. The tracks to be processed do not have to be allocated to a data set; you can select any tracks within the physical limitations of the input or output disk. Ranges of tracks to be processed are identified by a starting and ending track address (cylinder and head number).

With absolute track restores, the original tracks can be restored only to the original track addresses; you cannot restore to a new location on the output disk with absolute track addresses. You can do absolute track restores from any FDR, DSF, ABR, or SAR backup, including a DSF absolute track backup, as long as the requested tracks are on the backup data set. Absolute track restores must be like device restores, to the same device type (e.g., 3390 to 3390).

20.01 CONTINUED**DSF PRINT
OPTION**

DSF provides an option to print the contents of disk tracks by fully-qualified data set name, by absolute track address, or by using generic data set name selection. For each track selected, DSF will print the record zero (R0) plus each physical record on the track. The count field, key (if any) and data are printed in storage dump format (hexadecimal plus EBCDIC).

SMS SUPPORT

On a system with IBM's SMS (System Managed Storage) active, DSF supports SMS management of data sets.

When dumping, DSF (as well as FDR and ABR) will backup SMS information from the VVDS and VTOC for both VSAM and non-VSAM data sets. This includes SMS class information (storage, management, and data classes), and SMS indicators. For those data set types supported only on SMS-managed volumes (PDSEs, HFS data sets and Extended Format (EF) data sets), all information necessary to recreate the data set is also preserved.

When DSF restores a data set on a system with SMS active, SMS will be invoked for every data set which must be allocated, to decide if the data set should be managed by SMS, or allocated as non-SMS. The SMS storage class and management class ACS (Automatic Class Selection) routines will be invoked; they will be passed input class names:

- if the user specified STORCLAS=, NULLSTORCLAS, MGMTCLAS=, or NULLMGMTCLAS, those overriding values will be used.
- if the storage class or management class (or both) were not overridden by the user, the class associated with the input data set will be used.
- if the input data set was not SMS-managed, null classes will be passed.
- Your ACS routines may accept those classes, or override them with different values or even null values.

If SMS assigns a storage class to a data set, it will be SMS-managed; SMS will be invoked again to allocate the data set on a volume chosen by SMS. If no storage class is assigned, DSF will allocate the data set on a non-SMS volume (a target non-SMS volume must be indicated by a DISKx DD statement or a NVOL= operand on the SELECT statement). Even if data sets are allocated by SMS on a number of different volumes, DSF will restore those data sets in one pass of the backup file.

So, DSF can be used to convert data sets to SMS management, simply by updating the SMS ACS (automatic class selection) routines to assign storage classes to the restored data sets, or by specifying a storage class via the STORCLAS= operand on the SELECT statement. Data sets can be converted back to non-SMS if the ACS routines assign no storage class or the NULLSTORCLAS operand is specified. However, FDRCOPY ([Section 21](#)) or FDRCONVT ([Section 70](#)) may be a better choice for conversion of data sets.

Storage administrators, with proper authority, can override or bypass many of the SMS functions, to directly specify SMS classes, or to specify the volume serial on which SMS data sets are to be restored, by use of the BYPASSACS and BYPASSSMS operands on the RESTORE statement.

More detail on SMS support is in [Section 60](#).

20.01 CONTINUED

**PROCESSING
SPECIFIC
TYPES OF
DATA SETS**

Details of DSF and ABR processing for various types of data sets are in [Section 80.11](#).

**FDR
INSTANTBACK
UP**

If you are also licensed for FDR InstantBackup, it enhances DSF to provide the ability to take backups which are frozen at a given point-in-time. This works with almost any disk subsystem. It allows you to capture a point-in-time image of an online disk to an offline disk, effectively preserving the image of the online disk at that point-in-time. FDR InstantBackup can then read the offline disk and create the required backups or copies, without relabeling the disk or bringing it online. FDR InstantBackup for various hardware platforms is described in [Sections 25-29](#).

FDR InstantBackup also enables support for the HSDM (High Speed Data Mover) feature available on some disk subsystems. HSDM allows FDR to backup and restore data in an internal compressed format. If your disk subsystem includes HSDM, you can invoke HSDM support by adding the DCT=YES operand on DUMP statements. [See Section 80.33](#) for details.

20.02 DSF PROCESSING OPTIONS, PERFORMANCE AND REQUIREMENTS**DSF
OPERATIONS**

The first control statement in the DSF input must be a DUMP, PRINT, or RESTORE statement identifying the operation. It is followed by one or more SELECT statements identifying the actual data sets (or absolute track extents) to be processed. EXCLUDE statements are optionally used to exclude a subset of the selected data sets, e.g., all data sets starting with A except those starting with AB.

DSF will accept up to 250 SELECT/EXCLUDE statements in a single execution, unless that limit is overridden. Since a SELECT can select many data sets DSF can process any number of data sets in one execution.

**DSF BACKUP
OPTIONS**

A DUMP statement, with various operands to specify backup options, invokes a DSF backup. One or more SELECT statements follow it, identifying data to backup using one of 4 formats:

A. DATA SET NAME FORMAT

Dump data sets by identified by the DSN= operand, which can be a fully-qualified name, a dsname prefix, or a selection mask. The volumes identified by DISKx DD statements will be scanned for data sets and clusters matching the SELECT statements.

B. DD NAME FORMAT

Dump the data set specified on the DD Statement named by the DD= parameter. Note that DSF will get the dsname from the DD statement, but not any volume information. It will search the disk volumes identified by DISKx DD statements for the data set. This option can be used to select GDG generations by relative generation number.

C. DUMP ALL DATA SETS

Dump all data sets on the disk volumes specified by DISKx DD statements, using the ALLDSN operand.

D. ABSOLUTE TRACK ADDRESS FORMAT

Dump the tracks within the bounds specified by the FROM and TO address operands.

DSF will backup all selected data sets or tracks from all disk volumes specified by the DISKx DD statements to the tapes identified by corresponding TAPEx DD statements, in the order that the DISKx DDs appear in the JCL.

DSF also supports the ATTACH and MAXTASKS= operands, which will concurrently dump the disk volumes referred to by the DISKx DD statements to the tapes referred to by the corresponding TAPEx DD statements. Although there may be up to thirty-nine TAPEx/DISKx pairs, the TAPEx DDs must not specify more than nine unique tape units. The TAPEx DDs may include UNIT=AFF or VOL=REF to use the same drive for multiple backups or to place multiple backup files on the same tape volume; the number of unique units determines the number of concurrent DUMPs. FDR will serialize the backup of any disk volumes whose TAPEx DDs point to the same tape unit.

**DSF PRINT
OPTIONS**

A PRINT statement invokes a DSF track print. It is followed by one of more SELECT statements identifying the data to be printed using the same formats used for backup above. DSF will print the selected data from each volume identified by the DISKx DD statements to the print output data sets identified by the corresponding TAPEx DD statements.

20.02 CONTINUED

**DSF RESTORE
OPTIONS**

A RESTORE statement, with various operands to specify global restore options, invokes a DSF restore. One or more SELECT statements, identifying data to restore using one of 5 formats, follow it:

A. DATA SET NAME FORMAT

Restore the data sets or clusters specified by the DSN= operand to the data sets with their original names. DSN= can specify a fully-qualified name, a dsname prefix, or a dsname selection mask.

B. DATA SET NAME/NEWNAME FORMAT

Restore the data sets or clusters specified by the DSN= operand to new data set names, identified by the NEWNAME=, NEWGROUP=, or NEWINDEX= operand.

C. DD NAME FORMAT

Restore the data set specified on the DD statement named by the DD= operand to the data set of the same name.

D. DDNAME/NEWNAME FORMAT

Restore the data set specified on the DD statement named by the DD= operand to the data set named by the NEWNAME=, NEWGROUP=, or NEWINDEX= operand.

E. ABSOLUTE TRACK ADDRESS FORMAT

Restore the tracks within the bounds specified by the FROM and TO address operands.

If multiple restores are requested (multiple TAPEX/DISKx pairs), the restores are done serially in the order that the TAPEX DDs appear in the JCL unless MAXTASKS= is specified to request multiple concurrent restores. Similar to backup, FDR will recognize if certain TAPEX DD statements point to the same tape unit and will serialize restores from that unit.

DSF will search each backup data set identified by TAPEX DD statements for the data sets identified and will restore them. You can supply DISKx DD statements to identify the target volume for data sets restored from the equivalent TAPEX DD statements, or you can identify the target volumes via operands on the SELECT statements. Data tracks backed up by absolute track (FROM/TO) can only be restored by absolute track and must be restored to their original track addresses.

**MEMORY
REQUIRE-
MENTS**

The DSF DUMP basic memory requirement is identical to FDR's Memory Requirement. (See [Section 10.02](#)). However, the region may need to be increased if you have more than 250 control statements or more than 600 ICF VSAM or SMS data sets on a volume to be dumped.

DSF RESTORE processing requires a below-the-line region of 512K plus about 512 bytes for each data set or track range to be processed. ICF VSAM clusters may add an additional 1K bytes per component processed.

Some logical RESTORE operations may require additional memory, so a region of 1024K or more is recommended. You may want to specify REGION=0M to get the largest possible region. DSF will use only the storage it needs regardless of the REGION size.

**DSF ERROR
DETECTION**

If any of the selected data sets are not found on any input disk (dump) or backup (restore) or have errors (such as I/O errors), DSF will continue the dump or restore operation for the remainder of the data sets. An error message will identify the failing data set and a U0888 abend will be issued for the step at completion to call attention to the error.

20.02 CONTINUED

**COMPRESS
OPTION**

FDRDSF can be instructed to compress the data on the sequential backup file using Innovation's own proprietary software compression algorithm.

It is **not recommended** for backups to tape attached by ESCON or FICON channels because of the speed of the channel. Software compression will be ignored for backups created with the HSDM disk hardware option (DCT=YES) since the data is pre-compressed.

NOTE: all FDR restores will automatically recognize a compressed backup file and decompress it. No special option is required to restore a compressed backup.

**DUPLICATE
TAPE OPTION**

DSF has an option to create a duplicate or second copy of the backup tape during dump processing. When several volumes are dumped duplicate backup files may be made for one or more of the disks regardless of the others.

While dumping a disk to a TAPEX DD statement, the duplicate backup will be written to the TAPEX DD statement (same "x" value twice) if it is present. You may have TAPEX DDs for some TAPEX DDs and not for others in the same step.

Memory requirements do not increase with the use of the duplicate tape option.

**STORAGETEK
EXHPDM
SUPPORT**

FDRDSF supports the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM takes multiple concurrent tape outputs (such as FDR backup TAPEX or TAPEX DD statements) and directs them to a smaller number of tape drives, interleaving the data in a single tape file. ExHPDM is invoked by adding the SUBSYS= operand to the TAPE DD statements. [See Section 80.33](#) for more details.

SECURITY

Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system which supports the SAF (Security Authorization Facility) interface, such as ACF2 or TOP SECRET), you can enable the ALLCALL option. For DSF this results in these security checks:

- for data set backups and prints, DSF will always check to see if your userid has at least READ authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority, then DSF will check if you have at least READ authority under the DATASET security class to every data set being backed up or printed. Any data sets to which you are not authorized will be bypassed with an error message.
- for data set restores, DSF will check if you have at least UPDATE authority under the DATASET security class to every data set restored. Any data sets to which you are not authorized will be bypassed with an error message. If an output data set must be allocated, the operating system will check if you have CREATE/ALLOCATE authority for the data set (this is done even if ALLCALL is not enabled).

20.02 CONTINUED

**DATA SET
ENQUEUE
OPTION**

You can request, via the DSNENQ= operand, that each data set being dumped or restored be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name of SYSDSN and a minor name of the data set name itself, for each selected data set found in the VTOC of the input disk; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. Note that FDR cannot tell if the data set is being used for input or output. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations.

If you have requested data set ENQs, any data set that is in use will cause a FDR158 warning message to be printed; this will set the job error flag and will cause a U0888 abend when the step is complete (see "Step Termination" below). If you don't want in-use data sets to be considered an error, specify the ENQERR=NO operand; this prints the FDR158 message without setting the error flag.

Optionally you can request that inactive data sets be enqueued to DSF during the backup, to insure that no other job or TSO user can access the data set until the backup is done.

For backups, in-use data sets will still be dumped by default, but you must be aware that the backups of data sets which are being updated during the backup may be unusable, depending on the nature and format of the data. If you wish to bypass the backup of active data sets, specify the ENQERR=BYPASS operand.

For restores, DSF will attempt to enqueue any data sets that it allocates on the output disks, to insure that no other task tries to use them until the restore is complete, but if the ENQ fails, the data set is still restored. But for existing data sets, if the ENQ fails, the restore will be bypassed.

The DSNENQ= operand has 4 possible values:

USE -- data sets will be enqueued for the duration of the backup or copy from this disk volume.

For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option.

TEST -- data sets will only be tested to see if they are enqueued to another task at the time that the backup or copy from this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the dump is proceeding.

HAVE -- The data sets will be enqueued for the duration of the dump. If a data set is in use, the MVS operator must interact with FDR to decide how to proceed; a message (FDRW27) is issued to the MVS console, and the operator can respond:

WAIT - wait for the data set to become available; if it is not eventually dequeued, the FDR job may time out, so the operator must know which data sets are in use by long-running jobs or tasks.

NOWAIT - do not enqueue the data set. The FDR158 warning message is issued.

RETRY - try the enqueue again. If it fails again, the FDRW27 message is reissued.

NONE -- No data set ENQ will be issued. This is the default.

NOTE: If a data set name appears in a DD statement with DISP=SHR within the FDR job (not necessarily in the FDR step), and you specify DSNENQ=USE, HAVE or TEST, FDR will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the FDR job ends.

CAUTION: *This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.*

Use DSNENQ= to prevent other tasks from updating (or reading) data sets being dumped or restored. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

If HFS=QUIESCE is specified, special processing is done for HFS data sets (Hierarchical File System, used by OS/390 Unix System Services, OS/390 Unix). If the SYSDSN ENQ cannot be acquired, this may mean that the file system is mounted to OS/390 Unix, so FDR will attempt to quiesce the file system during the backup. Details on the quiesce function are found in [Section 80.11](#).

20.02 CONTINUED

**VTOC
ENQUEUE
OPTION**

FDR also supports, via the ENQ= operand, an ENQ on the VTOC of every volume being dumped, restored, or copied. For shared DASD, it can also invoke a hardware RESERVE on the volume during the FDR operation.

The VTOC is protected by an ENQ with major name SYSVTOC and a minor name of the volume serial. This ENQ is held by any task doing updates to the VTOC, including allocation of new data sets, extension of data sets to new extents, and scratching of existing data sets. This ENQ is normally of short duration, just for the few seconds necessary to update the VTOC, so if the ENQ is currently held by another task, FDR will wait for it to be released.

The SYSVTOC ENQ does not prevent access to existing data sets on the volume; it only insures that no other task is updating the VTOC while FDR is processing it. VTOC changes during a backup or copy could result in an invalid backup.

For disks shared with another MVS system or LPAR, ENQ=RESERVE requests that, in addition to the ENQ described above, a hardware RESERVE is done on the volume. RESERVE will prevent any system from reading or writing data on the volume, except for the system that FDR is running on, where only the ENQ protection applies. If you have a cross-CPU ENQ facility, such as GRS or MIM, you may be able to convert the RESERVE into a cross-CPU SYSVTOC ENQ and allow access to the volume during the operation (lookup SYSVTOC in the documentation for your product).

Use ENQ= to prevent other tasks from making changes to the VTOC during the backup or restore. Since DSF operates on individual data sets, there is usually no need to enqueue the VTOC as long as DSNENQ= (above) is specified. Member ENQ in the FDR ICL (Installation Control Library) has more information on VTOC ENQs

**STEP
TERMINATION**

If no errors occur during the execution of DSF, the DSF jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

- Some errors are critical. The jobstep ends immediately with a user ABEND.
- Some errors are critical only to a particular operation. For example, during a backup, some errors cause the backup of a particular disk to terminate immediately, but DSF may continue and attempt to backup other disks requested in the same step.
- Some errors are non-critical and the messages are warnings only. DSF will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with a U0888 abend to call your attention to the errors. Remember that a U0888 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer not to get a U0888 abend on a non-terminating error, the FDRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice ([see Section 90](#)).

20.03 DSF DUMP/PRINT JOB CONTROL REQUIREMENTS

To execute DSF for a DUMP or PRINT operation, the following JCL statements are required.

**STEPLIB or
JOB LIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRDSF resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRDSF), region requirement (REGION=, [see Section 20.02](#)), and optional PARM= operand.

If a PARM field is specified, DSF will use data specified as the first control statement, which must be a valid DUMP or PRINT statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRDSF,PARM='DUMP TYPE=DSF,DATA=USED'
//FDR EXEC PGM=FDRDSF,PARM='PRINT TYPE=DSF/ SELECT DSN=A.B.C'
```

If FDRDSF is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**SYSPRINT DD
STATEMENT**

Specifies the output message data set. It must be present and is usually a SYSOUT data set but it may be assigned to disk or tape. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**SYSPRINx DD
STATEMENT**

Specifies the output data set for messages related to the matching DISKx when the ATTACH option or MAXTASKS= option is used (see SYSPRINT for details). Must be present if ATTACH or MAXTASKS= is used but it is not used without them. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD.

**FDRSUMM DD
STATEMENT**

(Optional) if present, DSF will write one-line messages for each volume dumped, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set. FDRSUMM is used only if RTC=YES or DCT=YES is specified on the DUMP statement.

**SYSUDUMP DD
STATEMENT**

Specifies theabend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

20.03 CONTINUED

**DISKx DD
STATEMENT**

Specifies the input disk volume. The format will be:

```
//DISK1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial assigned of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). Only a single disk volume serial may be specified. You may use either DISP=OLD or DISP=SHR; it makes no difference.

The DISKx DD may point to a cataloged data set so that you can select that data set without knowing what volume it is on. For example,

```
//DISK1 DD DSN=MASTER.FILE,DISP=OLD
```

However, if the data set is cataloged as multi-volume, only the first volume in the catalog will be processed.

"x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character and must have a corresponding TAPEx statement, so there can be a maximum of 39 DISKx DDs. Processing will proceed for as many pairs of DISKx/TAPEx statements as are present, in the order that the DISKx DDs are present in the FDR JCL. If DUMMY is specified, this DD statement will be ignored.

For FDR InstantBackup

If you are also licensed for FDR InstantBackup, you can direct DSF to backup an offline point-in-time image of the volume to be backed up. This allows you to capture that point-in-time image and back it up at your leisure, even while updates are being done to the live, online disk. This may require special options on a DISKx DD statement. FDR InstantBackup is described in [Sections 25](#) through [29](#), with separate sections for each hardware platform on which FDR InstantBackup is supported.

20.03 CONTINUED

**TAPEx DD
STATEMENT**

Specifies the output data set for DUMP and the print data set for PRINT. "x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character. Multiple TAPEx DD statements may be present in the DSF step JCL; a unique value for 'x' must be used for each of them (e.g., TAPE1, TAPE2, etc.). There must be a TAPEx for each DISKx; if you code one or the other, not both, it will be ignored.

For DUMP Operations:

Specifies a tape or sequential disk data set to which the backup will be written. You must provide a TAPEx DD statement for each volume to be backed up in this step; the TAPEx DD will receive the backup of the volume specified by DISKx. If ATTACH is specified, FDR will attempt to attach that many concurrent backup subtasks, but may postpone some of them if it detects that they require a tape drive in use by another backup (see the notes on UNIT=AFF and VOL=REF below).

DUMMY is supported, for **testing** purposes only. The disk will be read, but the backup data will be discarded.

You must provide all the JCL parameters required to allocate and catalog the backup data set on disk or tape, which may include some or all of: DSN=, UNIT=, VOL=, SPACE=, and DISP=(NEW,CATLG). For tape, a volume count should be specified since the default is only 5 tape volumes, e.g., VOL=(, , , 255) .

DCB parameters are not required and should be omitted. However, tape unit hardware compaction (sometimes called IDRC), available on most tape cartridge drives) can be requested by adding DCB=TRTCH=COMP to your DD statement; Tape hardware compaction may be the default depending on local MVS options. For tapes attached by ESCON or FICON channels, Innovation recommends use of tape hardware compaction instead of FDR software compression (the COMPRESS= option).

```
Examples: //TAPE1 DD DSN=PROD.MVS001.BACKUP1,UNIT=TAPE,
//          VOL=( , , , 255 ),DISP=(NEW,CATLG)
//TAPE2 DD DSN=PROD.MVS002.BACKUP2,UNIT=DISK,
//          SPACE=(CYL,(500,100),RLSE),DISP=( ,CATLG)
```

For tape backups, UNIT=AFF or VOL=REF may be specified, referencing another TAPEx DD statement, to reduce the number of tape drives used in the step. UNIT=AFF=TAPEx will cause MVS to allocate the same tape drive for both DD statements, but will call for separate output tapes when each DD is opened. VOL=REF=*.TAPEx with LABEL=n can be used to stack multiple backup files on the same tape, providing more complete utilization of the tape volumes (which may be important as new technology increases tape volume capacity). FDR will automatically recognize that multiple TAPEx DDs point to the same tape drive and will serialize operations on that drive so that only one backup is directed to that drive at a time.

```
Examples: /* The following creates 2 backups on 2 different
//* tape volumes using the same tape drive. This may not
//* fully utilize the tape volumes but will allow for
//* concurrent restores from these backups.
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,UNIT=TAPE,
//          VOL=( , , , 255 ),DISP=(NEW,CATLG)
//TAPE2 DD DSN=PROD.MVS002.BACKUP2,UNIT=AFF=TAPE1,
//          VOL=( , , , 255 ),DISP=(NEW,CATLG)

//* The following creates a multi-file (and possibly
//* multi-volume) tape containing 3 backups.
//TAPE3 DD DSN=PROD.MVS003.BACKUP3,UNIT=TAPE,
//          VOL=( ,RETAIN , , 255 ),DISP=(NEW,CATLG)
//TAPE4 DD DSN=PROD.MVS004.BACKUP4,LABEL=2,
//          VOL=( ,RETAIN,REF=*.TAPE3( ,DISP=(NEW,CATLG)
//TAPE5 DD DSN=PROD.MVS005.BACKUP5,LABEL=3,
//          VOL=( ,RETAIN,REF=*.TAPE4( ,DISP=(NEW,CATLG)
```

20.03 CONTINUED

Your tape management software may require that you add an operand to the TAPEX DD to specify when the tape will be returned to the scratch pool. The operands are:

RETPD=nnnn or **LABEL=RETPD=nnnn** retain tape for "nnnn" days.

EXPDT=yyyyddd or **LABEL= EXPDT=yyyyddd** retain tape until Julian date yyyy.ddd

Certain expiration date values are treated as keywords by some tape management systems, e.g., EXPDT=99000 may indicate "expire when the backup data set is no longer cataloged". See [Section 80.32](#) for details.

If you are running the StorageTek ExHPDM (High Performance Data Mover) software product, you can direct FDR backups to ExHPDM with the SUBSYS= JCL operand, e.g.,

```
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,DISP=(NEW,CATLG),
// SUBSYS=(SOV,'CLASS(FDRBKUP)')
```

Please read [Section 80.33](#) and the ExHPDM program documentation for more details.

Backups can also be written to a sequential data set on disk, e.g.,

```
//TAPE1 DD DSN=TECH.BACKUP1,UNIT=3390,VOL=SER=TECH01,
// SPACE=(CYL,(25,5),RLSE),DISP=(,CATLG)
```

Although the backup will usually take less space than the original data being backed up, this can vary, so be sure to specify a secondary allocation quantity and the RLSE parameter to release unused cylinders.

For PRINT Operations:

Specifies the destination of the printed output from the PRINT statement. This is usually SYSOUT, but it may be directed to a disk or tape data set.

Example: //TAPE1 DD SYSOUT=*

**TAPExx DD
STATEMENT**

Specifies a second backup data set, using the same format documented for TAPEX (DUMP operation) above. A backup identical to TAPEX will be produced on TAPExx; the same data blocks are written to both simultaneously. For example, if DISK6 is being dumped to TAPE6, the inclusion of a TAPE66 DD statement will cause a second backup file to be produced.

TAPExx is optional; if omitted, no duplicate backup is created. TAPExx must be omitted for a PRINT or RESTORE operation.

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for DSF. Usually a DD * data set.

It is required, but if control statements were provided in the EXEC PARM=, it can be DUMMY.

20.04 DSF DUMP STATEMENT

DUMP D	TYPE=DSF	,ENQERR=BYPASS <u>PROCESS</u>
	,ATTACH MASTASKS=n	,FORMAT= <u>NEW</u> SPLIT
	,BUFNO=MAX nn	,HFS=QUIESCE
	,COMPRESS=ALL COPY1 COPY2	,ICFCORE=nnnnnn
	,DATA=ALL <u>USED</u>	,MAXCARDS=nnnnnn
	,DSNENQ= <u>NONE</u> TEST USE HAVE	,MAXERR=nnnn
	,DCT=YES <u>NO</u>	,RTC=YES <u>NO</u>
	,ENQ=ON <u>OFF</u> RESERVE	,SELTERR=NO YES
	,ENQERR=NO	,SNAP=(USE,REL)

DUMP STATEMENT The DUMP statement requests a dump (backup) operation. It must be the first statement input. Only one DUMP, RESTORE or PRINT statement is allowed per execution.

These operands will control the backup of each specified disk (DISKx DD statement) to its backup data set (TAPEx DD statement). If a TAPExx DD is also present, a second duplicate copy of the backup is written there. The backup data sets may be on tape or disk. The backups will proceed for each valid DISKx/TAPEx pair in the JCL. If the ATTACH or MAXTASKS= option is specified, a SYSPRINx DD is also required, to separate messages for each concurrent dump subtask.

One or more SELECT statements and optionally one or more EXCLUDE statements to specify the data sets or tracks to be dumped must follow it. However, if only the VTOC is to be dumped (DSN=VTOC is specified on the DUMP statement), no SELECT statements are required.

OPERANDS

TYPE=DSF	Specifies that a data set backup is to be performed. It is required.
ATTACH	Backups will proceed concurrently for all disks for which a triplet of DISKx, TAPEx, SYSPRINx DD statements appear. However, backups to TAPEx DD cards that specify the same tape drive will be automatically serialized. Up to 39 TAPEx DDs may be specified, but they must not point to more than 9 unique tape units (use UNIT=AFF= or VOL=REF= to reduce the number of allocated drives. Default: each DISKx/TAPEx pair will be processed one at a time, in the order that the DISKx DDs appear in the JCL, unless MAXTASKS= is specified.
BUFNO=	specifies how many buffers will be used for dumping each disk volume. Each buffer holds one disk track. The buffers acquired will be divided into 2 sets in order to overlap input and output I/O operations; each disk I/O will read disk tracks into one half of the buffers. MAX - buffers sufficient to read 1 cylinder of the input disk are acquired. nn - the specified number of buffers is acquired. Default: MAX. Innovation recommends that you do not override the default. However, BUFNO=2 will be forced when a backup (output) data set is on disk.

20.04 CONTINUED

- COMPRESS=** Controls the use of FDR software compression. Values for COMPRESS= are:
- ALL** -- the backup file for both copies (TAPEx and TAPExx) is to be compressed.
- COPY1** -- only the backup on TAPEx DD statements will be compressed.
- COPY2** -- only the backup on TAPExx DD statements will be compressed.
- See "Memory Requirements" in [Section 10.02](#) for the additional storage required by COMPRESS=. COMPRESS= is ignored if DCT=YES is also specified.
- Default: backups will not be compressed..
- COMPRESS is recommended for backups to disk files, and for tape backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of IDRC (tape hardware compression) usually results in better performance.
- DATA=** **USED** -- only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be backed up. On most volumes, this will make the dump run faster.
- ALL** -- all allocated tracks of all selected data sets will be backed up. You may need to specify DATA=ALL if the data sets to be backed up include JES2 spool data sets or CICS log data sets, since they usually do not have valid last block pointers.
- Default: USED.
- DCT=** DCT= is valid only if you are licensed for FDR InstantBackup. It will be honored only if the disk being backed up is in a disk subsystem with the HSDM option (High Speed Data Mover). HSDM allows FDR to backup and restore the internal compressed images of disk tracks, improving backup elapsed times up to 60%. It can also be specified as DUMPCOMPRESSEDTRACK=.
- YES** -- use HSDM for any volume where the disk hardware has the HSDM feature installed. Normal backup will be used for other volumes.
- NO** -- do not use HSDM.
- Default is NO.
- Note that DCT=YES implies RTC=YES; see the description of RTC= for its benefits.

20.04 CONTINUED

DSNENQ=

Specifies whether all of the data sets selected for backup will be ENQed. See "Data Set Enqueue Option" in [Section 20.02](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a warning message is issued for the data set but the data set will still be dumped unless the ENQERR=BYPASS operand is specified. A successful ENQ will prevent any other task from using the data set until the backup of that volume is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be dumped. The options for DSNENQ= are:

USE -- The data sets will be enqueued for the duration of the backup from this disk volume. This is the most frequently used option.

TEST -- The data sets will only be tested to see if they are enqueued to another task at the time that the dump from this volume starts.

HAVE -- The data sets will be enqueued for the duration of the dump. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not ENQ the data set)

RETRY (try the ENQ again)

NONE -- No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: NONE.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the backup is complete. However, use DSNENQ=NONE (or omit DSNENQ=) when another data set by the same name on another volume may be in use (e.g., restoring data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

ENQ=

Specifies whether an ENQ should be done on the VTOC of each disk volume while data sets from it are being backed up. See "VTOC Enqueue Option" in [Section 20.02](#) for more details.

ON -- the VTOC of each disk volume will be ENQed during its backup. This ENQ may be effective only on the system where the backup is executing; other systems may still be able to update the VTOC.

RESERVE - in addition to the ENQ, a hardware RESERVE will be issued on each disk volume during its backup. This is meaningful only on a system with "shared DASD" where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF -- the VTOC will not be enqueued or reserved during the backup.

Default is OFF.

20.04 CONTINUED

- ENQERR=** **NO** - If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see "Step Termination" in [Section 10.02](#)). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally.
- Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.
- ENQERR=** Specifies processing if the DSNENQ= option finds that a data set is in use (enqueued):
- BYPASS** - do not backup an active data set.
- PROCESS** - backup a data set even if it is active (a warning message will still be produced).
- Default: PROCESS.
- NOTE:** both ENQERR=NO and ENQERR=BYPASS/PROCESS may be specified on the same DUMP statement.
- FORMAT=** Specifies the format of the sequential backup file.
- NEW** -- the backup will be created using a maximum of a 56K blocksize. A block will contain the image of one or more tracks from the input disk.
- SPLIT** -- the backup will be created using a maximum blocksize of 32K. For blocks that would be more than 32K, they are written as 2 blocks of 32K or less. However, FORMAT=SPLIT causes use of a disk I/O technique which is less efficient than that used by FORMAT=NEW, which will impact backup performance.
- WARNING:** If you use a normal copy program (ex: IEBGENER) to copy a backup file created with FORMAT=NEW, you will not get any error messages, but the resulting tape will not be usable for a restore. Tapes in the new format must only be copied with the INNOVATION provided tape copy program (FDRTCOPY) or FATAR.
- Default: NEW if backup on tape -- SPLIT if backup is on disk.
- HFS=** **QUIESCE** invokes special processing when HFS (Hierarchical File System) data sets are backed up. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to (Unix System Services), so a "quiesce" call is issued to prevent updates to the HFS data set during the backup. If the quiesce fails and ENQERR=BYPASS was specified, the HFS data set will not be backed up. See "HFS" in [Section 80.11](#) for details and security requirements.
- Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for **all** data sets being backed up, not just HFS data sets.
- Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table (see [Section 90](#)). If you use the default or cannot run with superuser status, you should unmount the HFS file system before the backup to be sure of getting a usable backup.

20.04 CONTINUED

- ICFCORE=** Specifies that the size of the table used to store the ICF VSAM cluster and component names. The backup must save all of the component names and their associated clusters which exist on the current input disk volume, in order to match up VTOC DSCBs (with the component name) to cluster names (for selection). nnnnnn is specified in bytes and must be large enough to contain all the VSAM names (depending on the length and number of names).
- NOTE:** Specifying ICFCORE= will increase the backup region requirement by the value specified. The default value uses one of the dump buffers and imposes no additional memory requirement.
- Default: 53248, which will hold about 650 components.
- MAXCARDS=** Accept additional SELECT/EXCLUDE statements (over 250).
- Default: 250 statements.
- MAXERR=** The number of disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup is written to the ExHPDM subsystem, MAXERR=1 is recommended.
- WARNING: MAXERR over the default value may result in the loss of many data sets. This option should only be used when necessary and with care.**
- Default: 20 errors.
- MAXTASKS=** Specifies the maximum number of volumes that will be dumped concurrently in this step. It has the same affect as the ATTACH operand except that the maximum number of concurrent backups is limited to "n". The value may be from 1 to 9 but the actual number of concurrent backups will be no greater than the number of TAPEX DD statements in this step's JCL (if multiple TAPEX DD statements point to the same tape drive, only one concurrent dump will use that drive at any time). If MAXTASKS=2 or more, messages for each backup are directed to the SYSPRINx DD corresponding to the TAPEX DD associated with the backup ([see Section 10.03](#)).
- The default is 1 unless the ATTACH operand is specified.
- RTC=** It can also be specified as READTRACKCCW=.
- YES** -- use READ TRACK CCWs to read disk data tracks. RTC=YES also causes:
- up to 1 cylinder of disk data is read at a time.
 - FDR buffers are moved above the 16MB line (about 2MB per concurrent backup), allowing more concurrent backups to be run in one step.
 - the elapsed time of FDR backups when the backup data set is itself on disk is significantly improved.
- NO** -- use other CCWs to read disk data tracks.
- Default is NO.

20.04 CONTINUED

SELTERR=

Specifies what will happen at step termination if DSF finds that a SELECT or EXCLUDE statement was never referenced (no data set on any input disk was selected by the statement):

NO - a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.

YES -- a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.

Default: YES unless overridden in the FDR Global Option Table (See [Section 90](#)).

SNAP=

Used only when you are licensed for FDR InstantBackup and only if you have previously used program FDRSNAP to create an instant point-in-time image of a volume in an IBM RVA or StorageTek Iceberg/SVA disk subsystem with the Snapshot feature. [See Section 26](#) for details of the use of **SNAP=** with FDR.

USE - tells FDR InstantBackup to read the offline snapped copy of the online volume specified in JCL. FDR remembers the device address of the offline device most recently used as a target for a snap of each online volume unless an intervening IPL has occurred.

(USE,REL) - same as SNAP=USE, but at the end of backing up each snapped volume, FDR will issue a request to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the disk subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.

Default: backup the online disk for each volume

CONTINUED . . .

20.05 DSF PRINT STATEMENT

```

PRINT                                TYPE=DSF                                ,ENQ=ON|OFF|RESERVE
P                                    ,DATA=ALL|USED                            ,MAXCARDS=nnnnn
                                    ,DSN=VTOC

```

**PRINT
STATEMENT**

The PRINT statement requests a print operation. It must be the first control statement and must be followed by one or more SELECT statements and optionally one or more EXCLUDE statements to specify the data sets or tracks to be printed. However, if only the VTOC is to be printed (DSN=VTOC is specified on the PRINT statement), no SELECT statements are required.

DSF will print the data sets or tracks specified in storage dump format (hexadecimal and EBCDIC) and will write this printout to the TAPEX data set. The TAPEX print data is in print record format, with a record format of FB and logical record length of 121. This data set is usually SYSOUT or may be a sequential data set on tape or disk. You may specify the blocksize on the TAPEX DD statement as a multiple of 121 but the default blocksize on tape or disk is 1210.

OPERANDS

TYPE=DSF	Specifies that a data set print it to be performed. Must be specified on PRINT command.
DATA=	<p>USED -- only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be printed.</p> <p>ALL -- all allocated tracks of all selected data sets will be printed. Default: USED.</p>
DSN=VTOC	The VTOC and volume label track will be printed in addition to any selected data sets. If only the VTOC is required, you do not need any SELECT statements.
ENQ=	<p>Specifies whether an ENQ should be done on the VTOC of each disk volume while tracks are being printed. See "VTOC Enqueue Option" in Section 20.02 for more details.</p> <p>ON -- the VTOC of each disk volume will be ENQed during its print. This ENQ may be effective only on the system where the print is executing; other systems may still be able to update the VTOC.</p> <p>RESERVE - in addition to the ENQ, a hardware RESERVE will be issued on each disk volume during its print. This is meaningful only on a system with "shared DASD" where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.</p> <p>OFF -- the VTOC will not be enqueued or reserved during the print. Default: OFF.</p>
MAXCARDS=	<p>Accept additional SELECT/EXCLUDE statements (over 250). Default: 250 statements.</p>

20.06 DSF SELECT STATEMENT -- FOR DUMP OR PRINT

```

SELECT          DSN=filter
S
                DD=ddname

EXCLUDE        ALLDSN
X
                FROM(CYL=cccc,TRK=tttt),TO(CYL=cccc,TRK=tttt)
                ,DATA=ALL|NONE
                ,DSNENQ=NONE
                ,DSORG=(xx,xx,..)
                ,TAPEDD=x

```

**SELECT
STATEMENT
FOR DUMP or
PRINT**

This statement selects the data sets and/or tracks to be dumped or printed.

A **SELECT** statement identifies an individual data set, group of data sets, or range of tracks to be processed. An **EXCLUDE** statement identifies data sets from within those selected by **SELECT** statements which are not to be processed. All data sets in the VTOCs of DASD volumes specified by DISKx DD statements will be compared to these statements to identify those to be processed; each data set will be compared to each control statement until a match is found. It must match all criteria specified on the statement to qualify, e.g, if DSN= and DSORG= are both specified, it must be a data set with the right name, having the indicated organization. A maximum of 250 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since DSF will only dump or print data sets which are selected, EXCLUDE statements are required only to exclude certain data sets from within a larger group on a SELECT statement, and can also be used to EXCLUDE certain tracks.

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```

EXCLUDE DSN=A . B . **
SELECT  DSN=A . **

```

Example 2. Select all data sets except partitioned (PDSs):

```

EXCLUDE ALLDSN,DSORG=PO
SELECT  ALLDSN

```

**DUMPING ICF
VSAM
CLUSTERS**

ICF VSAM clusters can be selected by specifying the fully-qualified base cluster name or matching on the base cluster with generic data set selection. When selected, all components of that cluster that exist on the volumes being processed will be dumped or printed, including alternate indexes and key range components. DSF will not examine ICF VSAM component names when processing SELECT/EXCLUDE statements; components will be selected only if their cluster name is selected. For further information, [see Section 80.13](#), "VSAM Special Considerations".

20.06 CONTINUED

OPERANDS

DSN= Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80](#). This name or filter will be used when scanning the VTOCs of selected volumes.

EXAMPLES: DSN=USER1.JCL.CNTL

DSN=**LIST

DSN=PROD++.**.LIB*

DSN= cannot select GDGs by relative generation number; use DD= for that purpose.

DD= Specifies that a data set name is to be taken from a DD statement. This operand must point to the DDNAME of a JCL statement. Using this option enables the user to specify a non-standard data set name or a generation data set (GDG) relative generation.

EXAMPLE: SELECT DD=DD1

//DD1 DD DSN=A.B.C(0),DISP=SHR

Note that although DD= copies the data set name from the DD statement, it does not use the volume pointed to by that DD. The data set will be searched for only on the volumes pointed to by DISKx DD statements; the volume containing the data set must be among them. You could have a DISKx pointing to a data set and use DD=DISKx, so that the volume containing the data set will automatically be included.

ALLDSN Specifies that all the data sets on the disk volumes specified are to be selected. If the DSORG operand is specified, only those data sets with a matching organization are selected.

FROM/TO Identifies a range of track addresses for an absolute track dump or print. Absolute track statements can be mixed with SELECT statements for data sets. The FROM and TO operands must appear on the same input record, and cannot be continued. FROM/TO cannot be used on an EXCLUDE statement. The cylinder (CYL=) and track (TRK=) addresses are in decimal, relative to zero. Leading zeros can be omitted. For example, valid specification on a 3390-3, which has 3339 cylinders and 15 tracks/cylinder, are CYL=0 TRK=0 to CYL=3338 TRK=14. The FROM address must be lower than or equal to the TO address.

If the FROM track (TRK=) is omitted, zero is assumed. If the TO track (TRK=) is omitted, the last track of the cylinder is assumed. So, if you specify only cylinder numbers (e.g., FROM(CYL=5),TO(CYL=7)) all tracks in those cylinders are processed.

NOTE: DSN, DD, ALLDSN and FROM/TO are mutually exclusive. One and only one of these operands must be provided on each SELECT or EXCLUDE card.

20.06 CONTINUED

- DATA=** **ALL** -- all allocated tracks in the selected data sets will be dumped or printed. This can be used to override DATA=USED operand (the default) from the DUMP or PRINT statement for certain data sets.
- NONE** -- no data tracks will be backed up for any type of data sets selected. The backup data set will consist only of the FDR control records necessary to allocate and update the characteristics of the output data sets during a restore. A DATA=NONE backup can be used to allocate all the ddat sets on the backup on output volumes, as long as the data in those data sets will be recovered by other means (such as data base backups). Member DATANONE in the FDR ICL (Installation Control Library) has more information on the use of DATA=NONE.
- Default: only used tracks will be processed for PS and PO data sets unless DATA=ALL was specified on the DUMP or PRINT statement.
- DSNENQ=** **NONE** -- the data set enqueue will not be done for the selected data sets. This can be used to override the DSNENQ= operand on the DUMP/PRINT statement for certain data sets that you know will probably be ENQed by another task.
- Default: enqueue option is determined by the DSNENQ option specified on the DUMP or PRINT statement.
- DSORG=** Data sets are not to be selected by this SELECT unless their DSORG matches one of the DSORGs specified. If more than one DSORG is specified, they must be enclosed in parentheses.
- VALID DSORGs are:
- | | | |
|----------------|-------------------|-----------------|
| DA -- BDAM | PS -- SEQUENTIAL | UM -- UNMOVABLE |
| EF -- ICF VSAM | PO -- PARTITIONED | UN -- UNDEFINED |
| IS -- ISAM | | |
- TAPEDD=** **x** -- specifies a single character matching the "x" in a TAPEX DD statement. If this operand is specified, then this SELECT/EXCLUDE will only apply to data sets on the input disk volume specified by the DISKx DD statement. TAPEDD= might be used when multiple DISKx DD statements point to the same volume to select which data sets are to go to which backups.
- NOTE:** TAPEDD= should be used when the DISKx DD statements specify dsnames, using the catalog to determine the volumes on which the data sets reside. In that case, it may happen that more than one DISKx DD statement points to the same volume. If so, every requested data set that resides on the same volume will be dumped once for each DISKx DD statement pointing to that volume, unless TAPEDD is specified.

20.07 DSF RESTORE JOB CONTROL REQUIREMENTS

To execute DSF for a RESTORE operation, the following JCL statements are required.

**STEPLIB or
JOB LIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRDSF resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRDSF), region requirement (REGION=, [see Section 20.02](#)), and optional PARM= operand.

If a PARM field is specified, DSF will use data specified as the first control statement, which must be a valid DUMP or PRINT statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRDSF,PARM='RESTORE TYPE=DSF'
//FDR EXEC PGM=FDRDSF,PARM='RESTORE TYPE=DSF/ SELECT DSN=A.B.C'
```

If FDRDSF is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**STEP CAT or
JOB CAT DD
STATEMENT**

For non-VSAM data sets, when DSF must catalog non-VSAM data sets, they will be cataloged in that user catalog instead of in the system catalog with the matching alias. This might be useful when you are creating a test system, to catalog test copies of production data sets in a test catalog.

However, there is one important exception: if the data set being cataloged is a GDG generation, the STEP CAT/JOB CAT must contain a GDG base for that GDG; if not, it will ignore the STEP CAT/JOB CAT and catalog into the regular aliased catalog, possibly deleting other valid generations.

For VSAM clusters, the target catalog is controlled by the ICFCAT= operand, described in [Section 20.05](#). With the proper ICFCAT= option, the STEP CAT/JOB CAT may be honored.

STEP CAT/JOB CAT should not be used if any data sets being restored are SMS-managed.

**SYS PRINT DD
STATEMENT**

Specifies the output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD.

**SYS DUMP DD
STATEMENT**

Specifies the abend data set. Usually a SYSOUT data set. A SYS DUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

**TAPEx DD
STATEMENT**

Specifies an input backup data set on tape or disk, from which data is to be restored. "x" may be single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character. There may be up to 39 such TAPEx DD statements; DSF will restore from each in turn in the order they appear in the JCL. Restores are always done one backup at a time, there is no option for concurrent data set restores.

Each TAPEx DD must point to a FDR, DSF, ABR or SAR backup on tape or disk. Backups may not be concatenated on a single TAPEx; use multiple TAPEx DDs if multiple backups are to be restored.

If the backup to be restored was directed to the StorageTek ExHPDM software product as described earlier, you must specify the SUBSYS= operand on the TAPE DD statement to invoke ExHPDM and restore that backup, e.g.,

```
//TAPE1 DD DSN=PROD.MVS001.BACKUP1,DISP=OLD,SUBSYS=SOV
```

20.07 CONTINUED**DISKx DD
STATEMENT**

Specifies an output disk volume. The format will be:

```
//DISK1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial assigned of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). Only a single disk volume serial may be specified. You may use either DISP=OLD or DISP=SHR; it makes no difference.

"x" may specify any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character and must have a corresponding TAPEX statement. If DUMMY is specified, this DD statement will be ignored.

The DISKx DD statement is optional; DSF will dynamically allocate required output volumes if necessary. If present, the volume pointed to by the DISKx DD statement will be used to as the default target volume for restoring data sets selected from the TAPEX backup file, although this target volume may be overridden by SMS or the NVOL= operand. Complete rules for output volume selection are in [Section 20.08](#).

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for DSF. Usually a DD * data set.

It is required, but if control statements were provided in the EXEC PARM=, it can be DUMMY.

20.08 DSF RESTORE STATEMENT

RESTORE	TYPE=DSF	,MAXCARDS=nnnn
R		
	,BLKF=nn	,MAXERR=nnnn
	,BYPASSACS	,NOCAT
		,RECAT
	,BYPASSSMS	
		,PRESTAGE
	,CATIFALLOC	
		,RLSE
	,CONFMESS=<u>YES</u> NO	,%FREE=nn
	,DATA=ALL <u>USED</u>	,SELTERR=NO <u>YES</u>
	,DSNENQ=<u>NONE</u> TEST USE HAVE	,VRECAT
	,ICFCAT=ORIGINAL STEP CAT ALIAS	
	,MSGDGD=<u>DEFERRED</u> ACTIVE ROLLEDOFF INPUT	

**RESTORE
STATEMENT**

The RESTORE statement is required for a RESTORE operation. Only one is allowed per execution, and it must be the first control statement. One or more SELECT statements must follow it, with optional EXCLUDE statements, to specify the data sets or tracks to be restored.

DSF will read each backup file specified on a TAPEX DD statement and will restore one or more data sets as indicated by the SELECT and EXCLUDE statements that follow (See [Section 20.09](#)). The backup files are read one at a time (in the order that the TAPEX DDs exist in the step JCL) but while reading a backup file, DSF can restore the selected data sets to one or more output disk volumes concurrently. The target disk volume will be selected for each data set by the following rules, in this order:

- If the NVOL= operand was specified on the SELECT statement which selected this data set, that volume or volumes will be used. See the description of NVOL= in [Section 20.09](#) for details of target volume selection.
- If the step JCL contains a DISKx DD statement matching the TAPEX DD from which the data set is being restored, then the disk to which it points will be selected for output.
- If the output data set name is cataloged, then the volume to which it is cataloged will be chosen. The output data set name will be the original name or the new name if a NEWNAME=, NEWGROUP=, or NEWINDEX= operand was specified on the SELECT statement which selected the data set. If the data set is cataloged as being on multiple volume serials, then the volser will be selected from that list based on the volume sequence number in the F1 DSCB (field DS1VOLSQ) of the input data set.
- If none of the above apply, then the serial of the volume from which the data set was dumped, as recorded on the backup data set, will be used.

If SMS (System Managed Storage) is active on this system, and the data set does not already exist on the volume selected by the rules above, SMS is invoked to decide if the data set should be SMS-managed. If so, SMS will select an output volume which may be totally different from the target volume chosen by DSF. SMS rules are detailed in [Section 20.01](#) and [Section 70](#).

20.08 CONTINUED

OPERANDS

TYPE=DSF

Specifies a data set restore. Must be specified on the RESTORE command.

BLKF=

Changes the blocksize (BLKSIZE) of certain fixed (RECFM=FB) and variable (RECFM=VB) data sets. BLKF= specifies the number of blocks per track (1 to 10) and is used to calculate a new blocksize based on the output track size (BLKF=1 results in BLKSIZE=32760). For RECFM=FB data sets the blocksize will be rounded down to a multiple of the logical record length (LRECL).

DSORG=PS (sequential) data sets will be physically reblocked during the restore; the new blocksize must be larger than the original blocksize except when restoring from 3390 to 3380.

DSORG=PO (partitioned) data sets will not be reblocked but the new blocksize will be used for new members; it must always be larger than the original blocksize.

Other data set types, including UN (undefined), DA (direct access), PSE (enhanced format sequential), POE (PDSE), HFS and all types of VSAM, will not be reblocked.

Default: data sets are not reblocked during restore; all original blocks will be restored without change, although they may be written to new locations. BLKF= is usually used when restoring to an unlike device type (e.g., 3380 to 3390) to improve utilization of the new track size but can also be used during like device restores.

BYPASSACS

On a system with SMS (System Managed Storage) active, the SMS ACS (Automatic Class Selection) routines are not to be invoked for data sets which must be allocated. If a data set has a SMS storage class assigned (see STORCLAS= in [Section 20.09](#)) it will be SMS-managed, and SMS will be invoked to allocate the data set on an SMS-chosen volume, but SMS will not be allowed to override the storage class or management class assigned to the data set.

Default: on an SMS system, the SMS ACS routines will be invoked for every data set which has to be allocated. The assigned storage and management classes will be passed to those routines, which can approve or override them. A data set will be passed to SMS for allocation if the storage class ACS routine assigns a storage class to the data set.

20.08 CONTINUED

BYPASSSMS

On a system with SMS (System Managed Storage) active, SMS data sets will be directly allocated on SMS-managed volumes, bypassing normal SMS storage group and volume selection. The selected output volume must be a SMS-managed disk volume, and the data sets being restored must have a SMS storage class assigned (see BYPASSACS above and STORCLAS= in [Section 20.09](#)). The data sets will be allocated and cataloged according to SMS standards.

Normal SMS facilities do not allow allocation of data sets on specific volume serials, but BYPASSSMS will do so, allowing data sets to be located for performance or other reasons. Note that if BYPASSACS is also specified, the assigned SMS classes will not be validity- or authority-checked.

Default: on an SMS system, for data sets which are SMS-managed and must be allocated, the SMS storage group ACS routine will be invoked to select a storage group and SMS will select a SMS-managed volume and allocate and catalog the data sets.

BYPASSACS and BYPASSSMS are primarily for use by storage administration personnel, since they bypass normal SMS allocation controls and rules. In order to use BYPASSACS or BYPASSSMS, the user of DSF must be authorized to the RACF profile

STGADMIN.ADR.RESTORE.BYPASSACS

in class FACILITY, or the equivalent in other security systems.

CATIFALLOC

Specifies that non-VSAM output data sets will be cataloged even if they were preallocated (not allocated by the restore); the output data set will be cataloged if it is not already cataloged on another volume (unless the RECAT operand was specified).

Default: output data sets are cataloged only when the restore allocates them.

CONFMESS=

YES -- before beginning any restore that includes an absolute track request (SELECT FROM/TO), FDR will request confirmation via a WTOR message (FDRW02) to which the MVS operator must reply.

NO -- suppress the WTOR and begins the restore immediately.

Default: YES.

DATA=

ALL -- all of the allocated tracks in each data set will be restored. ALL should not be specified unless the backup was taken using the DATA=ALL option. DATA=ALL should not be specified with the RLSE and %FREE operands.

USED -- only the used portion of the PS (physical sequential) and PO (partitioned) data sets will be restored.

Default: USED.

20.08 CONTINUED

DSNENQ=

Specifies whether all of the data sets being restored will be ENQed. See "Data Set Enqueue Option" in [Section 20.02](#) for more details.

If you are restoring over an existing data set and the ENQ fails, the restore will be bypassed with an error message. If the restore must allocate the output data set and the ENQ fails, no error message is issued and the restore is still done. A successful ENQ will prevent any other task from using the data set until the restore from the current backup data set is complete. An ENQ failure is considered an error but it will not prevent other data sets from being restored. The options for DSNENQ= are:

USE -- The data sets will be enqueued for the duration of the restore from the current backup data set. This is the most frequently used option.

TEST -- The data sets will only be tested to see if they are enqueued to another task at the time the restore starts. The data set will not be enqueued and other tasks may enqueue it while the restore is preceding.

HAVE -- The data sets will be enqueued for the duration of the restore. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not enqueue the data set)

RETRY (try the enqueue again; may result in the FDRW27 message again)

NONE -- No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: NONE. Note that NONE or TEST may allow other jobs to attempt to read the data set being restored before DSF has restored all of the data tracks. Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the restore is complete. However, use DSNENQ=NONE (or omit DSNENQ=) when another data set by the same name on another volume may be in use (e.g., restoring data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

20.08 CONTINUED

ICFCAT=

Applies to ICF VSAM files only. Specifies the source of the catalog name to be used if an output ICF VSAM cluster must be allocated.

ORIGINAL -- use the catalog in which the original dumped cluster was cataloged. When restoring a cluster to a new name, ICFCAT=ORIGINAL is treated like ICFCAT=ALIAS, described below. If you need to catalog the output cluster into the same catalog as the input cluster but that catalog is not the one aliased for the new name, you must specify ICFCAT=STPCAT and supply a STEPCAT DD statement pointing to that catalog.

STPCAT -- use the STEPCAT as the target catalog. If a STEPCAT DD statement is not supplied, it will use the master catalog or the catalog which is aliased for this data set in the master catalog.

ALIAS -- determine the catalog from the alias name in the master catalog. If no alias is found and the cluster is being restored to the same name, use the input cluster's original catalog. If no alias is found, and the cluster is being restored to a new name, DSF will use the STEPCAT (if present in the JCL) or the master catalog. Multi-level alias (MLA) is supported.

Default: ORIGINAL, except that if the cluster is being restored to a newname (NEWGROUP or NEWINDEX specified) the default is ALIAS. If the output cluster is SMS-managed, ALIAS is forced.

MAXCARDS=

Accept additional SELECT/EXCLUDE statements (over 250).

Default: 250 statements.

MAXERR=

Specifies the number of disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup was written to the ExHPDM subsystem, MAXERR=1 is recommended.

Default: 20 errors.

**NOCAT
RECAT**

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.

NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified.

20.08 CONTINUED

PRESTAGE	<p>Output data sets which already exist on the target output volume will not be restored. This may be used to avoid restoring data sets which have already been restored. If the output data sets do not exist on the target volume, they will be allocated and restored.</p> <p>Default: pre-allocated data sets will have their contents overlaid.</p>
RLSE %FREE=	<p>RLSE -- all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn -- a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement).</p>
SELTERR=	<p>Specifies what will happen at step termination if one or the SELECT or EXCLUDE statements was never referenced (no data set on any input disk was selected by the statement):</p> <p>NO - a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.</p> <p>YES -- a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.</p> <p>Default: YES unless overridden in the FDR Global Option Table (See Section 90).</p>
SMSGDG=	<p>Specifies the status of SMS-managed GDG (Generation Data Group) data sets, if allocated by the restore.</p> <p>DEFERRED, ACTIVE, or ROLLEDOFF will set the GDG to that status.</p> <p>INPUT will set the GDG to the original SMS status of the GDG generation, as recorded on the backup tape. If the original GDG was non-SMS, it will be set ACTIVE if that generation is currently cataloged, otherwise DEFERRED.</p> <p>Default: DEFERRED.</p>

20.08 CONTINUED

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- The required components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will restore on top of that existing cluster and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume).

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for some reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

20.09 DSF SELECT STATEMENT -- FOR RESTORES

SELECT	DSN=filter	,NOCAT
S	DD=ddname	,RECAT
	ALLDSN	
EXCLUDE	FROM(CYL=cccc,TRK=tttt),TO(CYL=cccc,TRK=tttt)	,NVOL=(vvvvvv,vvvvvv,...)
X		
	,BLKF=nn	,PRESTAGE
	,DATA=ALL NONE	,RLSE
		,%FREE=nn
	,DATACLAS=dataclass	,STORCLAS=storageclass
	,NULLDATACLAS	,NULLSTORCLAS
	,DSNENQ=NONE	
		,TAPEDD=x
	,MGMTCLAS=managementclass	
	,NULLMGMTCLAS	,TRK=nnnnn
		,CYL=nnnnn
	,NEWNAME=newname	
	,NEWGROUP=newgroup	,VRECAT
	,NEWINDEX=newindex	
	,NEWDD=newddname	

**SELECT
STATEMENT
FOR RESTORE**

This statement selects the data sets and/or tracks to be restored.

A SELECT statement identifies an individual data set, group of data sets, or range of tracks to be processed. An EXCLUDE statement identifies data sets from within those selected by SELECT statements which are not to be processed. All data sets in the backup files specified by TAPEX DD statements will be compared to these control statements to identify those to be restored; each data set will be compared to each control statement until a match is found. EXCLUDE statements should only contain the operands DSN=, DD=, ALLDSN, FROM, and TAPEDD=. A maximum of 250 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since DSF will only restore data sets which are selected, EXCLUDE statements are required only to exclude certain data sets from within a larger group on a SELECT statement, and can also be used to EXCLUDE certain tracks.

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A.B.**
SELECT  DSN=A.**
```

Example 2: Select all data set except those beginning with "SYS" on TAPE2:

```
EXCLUDE DSN=SYS**,TAPEDD=2
SELECT  ALLDSN
```

**NEWNAME/
NEWGROUP/
NEWINDEX for
ICF VSAM**

If you are restoring to a pre-allocated ICF VSAM cluster, where the cluster name is the same as the original, but the components are named differently, you must specify the cluster name as both DSN= and NEWNAME=; this causes DSF to LOCATE the new component names. NEWNAME cannot be specified for clusters with more than one alternate index.

If restoring an ICF VSAM cluster to a new name and the new cluster must be allocated, you should specify NEWGROUP= or NEWINDEX=. DSF will modify both the cluster and component names. This is illustrated in [Section 20.11](#).

ICF VSAM catalogs and VVDSs cannot be restored to a NEWNAME. DSF will properly restore an ICF VSAM catalog even if it has been re-allocated thereby generating a new CATINDEX name.

20.09 CONTINUED

See [Section 80.13](#) "VSAM Special Considerations".

OPERANDS**DSN=**

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the names of data sets on the backup tapes to be restored. For ICF VSAM clusters, only the cluster name is compared; you cannot select by component name.

```
EXAMPLES: DSN=USER1.JCL.CNTL
          DSN=**LIST
          DSN=PROD++.**.LIB*
```

You cannot use DSN= to select generation data group (GDG) generations by relative generation number. Use DD= if relative GDG numbers are required.

DD=

Specifies that a data set name is to be taken from a DD statement. This operand must point to the DDNAME of a JCL statement. Using this option enables the user to specify a non-standard data set name or a generation data set (GDG) relative generation.

```
EXAMPLE:  SELECT      DD=DD1
           //DD1      DD   DSN=A.B.C(0),DISP=SHR
```

ALLDSN

Specifies that all of the data sets dumped on the backup data set will be restored.

20.09 CONTINUED

FROM/TO

Identifies a range of track addresses for an absolute track restore. Absolute track statements can be mixed with SELECT statements for data sets. The FROM and TO operands must appear on the same input record, and cannot be continued. FROM/TO cannot be used on an EXCLUDE statement.

The cylinder (CYL=) and track (TRK=) addresses are in decimal, relative to zero. Leading zeros can be omitted. For example, valid specification on a 3390-3, which has 3339 cylinders and 15 tracks/cylinder, are CYL=0 TRK=0 to CYL=3338 TRK=14. The FROM address must be lower than or equal to the TO address.

If the FROM track (TRK=) is omitted, zero is assumed. If the TO track (TRK=) is omitted, the last track of the cylinder is assumed. So, if you specify only cylinder numbers, e.g. "FROM(CYL=5),TO(CYL=7)", all tracks in those cylinders are processed.

Absolute track restores are normally used with backups created by absolute track dumps, so that you know which tracks are on the backup. However, you can do absolute track restores from data set or full-volume backups if you know the track addresses available on the backup (a COMPAKTOR MAP or FDRABRP PRINT VTOC can identify those tracks).

CAUTION: On an absolute track restore, the requested tracks will always be restored to the same absolute address they were dumped from, regardless of what data set, if any, is allocated there. The VTOC will not be updated.

When selecting an output volume for a absolute track restore, the rules defined in [Section 20.08](#) are used except for cataloged output data set name (since there is none); normally a DISKx DD or NVOL= operand are used to identify the output volume. If the output volume is not the original volume, use FROM/TO with care to avoid overlaying valid data sets.

NOTE: DSN, DD, ALLDSN and FROM are mutually exclusive. One and only one of these operands must be specified on each SELECT statement.

ICF VSAM NOTE: The 'SYS1.VVDS' data set will be not be restored even if selected using the ALLDSN or DSN=filter operands. [See Section 80.14](#) for further details.

BLKF=

Selected PS and PO data sets are to be reblocked during the restore; see BLKF= in [Section 20.08](#) for details.

Default: data sets are not reblocked unless BLKF= was specified on the RESTORE statement. The restore will fail if the input data set has blocks larger than the track size of the output disk.

DATA=

ALL -- restore all allocated tracks in each data set selected by this SELECT statement.

NONE - do not restore any data tracks for data sets selected by this SELECT statement.

See DATA= in [Section 20.08](#) for details.

Default: restore only the used tracks of PS, PO and VSAM data sets, unless DATA= was specified on the RESTORE statement.

20.09 CONTINUED

DATACLAS=
NULLDATACLAS

On a system with SMS active, specifies the SMS data class to be associated with the data set being restored, overriding the original data class of the data set (if any). The Data Class ACS routine will not be invoked.

NULLDATACLAS changes the data class to null (not specified).

Default: the original data class of the input data set (if any) will be associated with the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is set.

DSNENQ=

NONE -- bypass the data set enqueue for data sets selected by this SELECT statement.

Default: the enqueue option is determined by the DSNENQ= operand specified on the RESTORE statement.

MGMTCLAS=
NULLMGMTCLAS

On a system with SMS active, specifies the SMS management class to be presented to the SMS Management Class ACS routine for the data set being restored, overriding the original management class of the data set (if any). The ACS routine may accept or override this class.

NULLMGMTCLAS changes the management class to null (not specified).

Default: the original management class of the input data set (if any) will be passed to the ACS routine for the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is passed.

NEWNAME=
NEWN=

Restore the selected data set with a new name. NEWNAME should only be used with DSN= or DD=, and should not be used for ICF VSAM clusters unless they are preallocated. If the newname ends in a GDG relative generation number, e.g., NEWNAME=gdgname(-1), a LOCATE will be done to get the proper absolute generation number.

NEWGROUP=
NEWG=

Restore the selected data sets using a new group name. The characters specified will replace the beginning of the input data set name. Care should be taken when periods are used that index levels are not incorrectly changed. DSF will check the new names for valid IBM standards.

EXAMPLE: SELECT DSN=ABC**, NEWG=XYZ

Any data sets restored will be renamed to start with characters XYZ.

20.09 CONTINUED

**NEWINDEX=
NEWI=**

Restore the selected data sets using a new name formed by adding or replacing one or more index levels in the original name; replacement index levels do not have to be the same length as the original indexes they replace. In the simplest case, each index level specified in NEWI is used in place of the corresponding index in the original name. Any remaining index levels at the end of the name are copied unchanged. This can easily be used to change the first indexes of the name.

For example, if the input data set is A.B.C.D,

NEWI=D	results in	D . B . C . D	(first index replaced)
NEWI=DD . E	results in	DD . E . C . D	(first 2 indexes replaced)

If a period is specified without any preceding characters, one original index level is copied from the input data set name to the output. This allows you to easily modify indexes in the middle of the name.

For example, if the input data set is A.B.C.D,

NEWI=. . E	results in	A . B . E . D	(third index replaced)
NEWI=FF . . . G	results in	FF . B . C . G	(first and fourth indexes replaced)

If + is specified before a new index level, that new index is inserted into the output data set name at that point. If ++ precedes the new index, it will be added to the end of the name. If - is specified, the next input index level will be dropped from (not copied to) the output name.

For example, if the input data set is A.B.C.D,

NEWI=+F	results in	F . A . B . C . D	(new first index added)
NEWI=. . +F	results in	A . B . F . C . D	(new third index added)
NEWI=++F	results in	A . B . C . D . F	(new last index added)
NEWI=. . -	results in	A . B . D	(third index dropped)
NEWI=Q . - . +E	results in	Q . C . E . D	(combination)

Note that, except for the ++ option, every period in the NEWI= mask corresponds to a period (one index level) in the original (input) data set name. The resulting new name will be checked to insure it meets IBM standards.

If the NEWI= value ends in a GDG relative generation number, e.g., NEWI=..NEWMAST(-2), that relative number will be added to the end of the newname, and a LOCATE done to get the proper absolute generation number.

NEWI= is a convenient way to rename every input data set, while using some index levels from the original name and replacing other indexes or adding new indexes.

NEWDD=

Specifies the name of a DD statement from which the new name of the output data set is obtained.

NOTE: NEWN=, NEWG=, NEWI=, and NEWDD= are mutually exclusive. If none of them are specified, the data set is restored under its original name. NEWN= and NEWDD= should not be used on SELECT statements which select more than one data set.

20.09 CONTINUED

NOCAT
RECAT

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged, unless overridden by NOCAT/RECAT on the RESTORE statement.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.

NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified on the RESTORE statement.

20.09 CONTINUED

NVOL=

Specifies the volume serial(s) of output disk volumes to which data sets selected by this statement are to be restored. You may specify:

- 1) A single specific volume serial, e.g.,
NVOL=ABRC123
- 2) A list of specific volume serials enclosed in parentheses,
e.g., NVOL=(TSO001,TSO002,TSO003)
- 3) A group of volumes by placing an asterisk at the end of the volser prefix,
e.g., NVOL=TSO*
- 4) A combination of specific and group, e.g., NVOL=(TSO*,PROD*,ABC001)
- 5) All online disk volumes may be selected by NVOL=*

A list of online target volumes matching your specification is generated by scanning all disk UCBs in the system UCB chains; there is no guarantee of the order in which UCBs are found, so you cannot predict the order of the volume serials in the list. If you specify volume serials or groups which are not online, they are ignored and no error message will result.

However, if the first or only specification is specific volume serial, it will be chosen as the first target volume, with other volumes placed after it in UCB chain order.

Also, if you are restoring a multi-volume data set (non-VSAM or SMS-managed VSAM), the volume sequence number of the piece of the data set being allocated will be used to select a specification from your list. For example, if NVOL=(A,B,C), the second piece of the data set will go to volume B. If that specification is a group, the first volume in the UCB chain matching that group will be tried. If the allocation is unsuccessful (such as insufficient free space), then other volumes in the NVOL list will be tried as described above.

The first target volume is checked to see if an output data set already exists there. If so, it restores over the existing allocation (unless PRESTAGE was specified). If not, it attempts to allocate the output data set on that volume. If the allocation fails, it will be retried on successive volumes in the list until it succeeds or until it fails on 64 volumes. If the list contains several disk device types, "like" volumes (same type as the data set being restored) will be tried first, then unlike devices.

For multi-volume data sets, a target volume is bypassed if a piece of the data set already exists there but is not the right piece, so that it will not attempt to restore the third volume of a data set on top of the first volume. When it finds a target volume in the list that does not contain a piece of the data set, it will be allocated.

Specifying multiple volsers or a volume group allows you to restore data sets in one pass even when no one volume has available space to contain them all; they will be spread across many of the target volumes.

Default: the output volume will be selected by rules defined in [Section 20.08](#). Note that when NVOL= is specified, and data sets are selected which are currently allocated and cataloged, DSF will restore them to the new volumes, and not to the volume on which they are cataloged.

On a system with SMS active, NVOL= may be ignored if the data set does not exist on the volume specified and the data set is SMS-managed (see STORCLAS= below).

Note: if an allocation is attempted on several volumes from your NVOL list but it fails on all of them, the message printed will usually show the allocation failure codes from the **first** volume only; failure codes from other volumes are not displayed and may be different.

20.09 CONTINUED

PRESTAGE

Specifies that selected data sets will not be restored if the output data set already exists on the first target output volume. This may be used to avoid restoring data sets which have already been restored.

Default: restore pre-allocated data sets, overlaying the existing contents of those data sets, unless PRESTAGE was specified on the RESTORE statement.

**RLSE
%FREE=**

RLSE -- specifies that all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.

%FREE=nn -- specifies a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.

Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation.

Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement or by RLSE/%FREE= on the RESTORE statement).

**STORCLAS=
NULLSTORCLAS**

On a system with SMS active, specifies the SMS storage class to be presented to the SMS Storage Class ACS routine for the data set being restored, overriding the original storage class of the data set (if any). The ACS routine may accept or override this class.

NULLSTORCLAS changes the storage class to null (not specified).

Default: the original storage class of the input data set (if any) will be passed to the ACS routine for the output data set. For a non-SMS input data set, a null class is passed.

If the Storage Class ACS routine assigns a storage class to this data set or accepts the class passed, the data set will be allocated as SMS-managed., and the SMS Storage Group ACS routine may be invoked to determine the actual target volume. If the Storage Class ACS routine returns a null (blank) storage class name, the data set will be allocated as non-SMS and the DSF rules listed in [Section 20.08](#) for volume selection must select a non-SMS target volume.

TAPEDD=

x -- specifies a single character matching the "x" in a TAPEX DD statement. If this operand is specified, this statement will apply only to data sets in the backup file pointed to by TAPEX.

Default: data sets specified will be restored from any or all of the backup files specified by the TAPEX DD statements on which they are found.

**TRK=
CYL=**

If the data set selected by this SELECT statement must be allocated, CYL= or TRK= specifies the number of cylinders or tracks to be allocated to the data set. On PS or PO files when DATA=ALL was not specified, this value should be at least equal to the used portion of the data set. On all other types of files and when DATA=ALL is specified, this value should be equal to or greater than the original size of the file. For ICF VSAM clusters, modifies the size of the base data component only. If the space is too small for the data being restored, the restore will automatically extend the file for non-VSAM.

Default: use the original size of the data set.

20.09 CONTINUED

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume) unless VRECAT was specified on the RESTORE statement.

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for any reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

20.10 DSF DUMP EXAMPLES

**DUMP
INDIVIDUAL
DATA SETS**

Dump specific data sets using the data set ENQ Option. The backup will be expired by tape management in 20 days. Four data sets are dumped from one disk volume to 3490E tape cartridges, referencing two by data set name and one through a DD statement. The fourth is an ICF VSAM cluster. DSF will enqueue the data sets for the duration of the dump. Since the tapes are attached by relatively slow parallel (bus/tag) channels, the backup will be compressed by FDR software compression.

```
//DUMP          EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=SYSDA,DISP=OLD,VOL=SER=RVS370
//TAPE1         DD     UNIT=3490,DSN=MYTAPE,DISP=(,KEEP),
//              LABEL=RETPD=20
//MASTER       DD     DSN=IAM.FILB(0),DISP=SHR
//SYSIN         DD     *
                DUMP    TYPE=DSF,DSNENQ=USE,COMPRESS=ALL
                SELECT  DSN=SYS2.LINKLIB
                SELECT  DSN=TSO.EDIT
                SELECT  DD=MASTER
                SELECT  DSN=VSAM.CLUSTER
/*
```

**DUMP ALL
DATA SETS**

Dump all the data sets, except system data sets, on two disk volumes, one volume at a time. ENQ and RESERVE the VTOCS of the volumes during the backup. Only the used tracks within PS and PO data sets will be dumped. Volume counts are specified on the TAPEX DD statements in case either backup requires more than the IBM default of 5 volumes.

```
//DUMPALL      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=3380,VOL=SER=RES001,DISP=SHR
//TAPE1         DD     UNIT=TAPE,DSN=DSF.BACKUP.A,DISP=(,KEEP),
//              VOL=(,99)
//DISK2         DD     UNIT=3390,VOL=SER=SCR001,DISP=SHR
//TAPE2         DD     UNIT=TAPE,DSN=DSF.BACKUP.B,DISP=(,KEEP),
//              VOL=(,99)
//SYSIN         DD     *
                DUMP    TYPE=DSF,ENQ=RESERVE,COMPRESS=ALL
                EXCLUDE DSN=SYS1.**
                SELECT  ALLDSN
/*
```

**ABSOLUTE
TRACK DUMP**

Dump four explicitly specified ranges of tracks. The first SELECT dumps cylinder 30 in its entirety (TRK= defaulted in both FROM and TO). The second SELECT will dump to the end of cylinder 45. All cylinder and track numbers are in decimal relative to zero.

```
//DUMPTRK      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISKW         DD     UNIT=3380,DISP=OLD,VOL=SER=A17865
//TAPEW         DD     UNIT=TAPE,DSN=MASTER,DISP=(,KEEP)
//SYSIN         DD     *
                DUMP    TYPE=DSF
                SELECT  FROM(CYL=30),TO(CYL=30)
                SELECT  FROM(CYL=42,TRK=5),TO(CYL=45)
                SELECT  FROM(CYL=50),TO(CYL=55,TRK=7)
                SELECT  FROM(CYL=66,TRK=8),TO(CYL=402,TRK=11)
/*
```

20.10 CONTINUED

**DUMP FROM
MULTIPLE
VOLUMES
CONCUR-
RENTLY WITH
EXHPDM**

Backup data sets from 3 disk volumes concurrently using the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM is invoked by the SUBSYS= operands on the TAPEx DD statements; see [Section 80.33](#) and the ExHPDM program documentation for details on the values to provide. These 3 concurrent backups will be combined into one file on a tape managed by ExHPDM.

```
//DUMP          EXEC  PGM=FDRDSF,REGION=0M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1        DD    UNIT=SYSDA,DISP=OLD,VOL=SER=111111
//TAPE1        DD    DSN=BACKUP.V111111,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN1     DD    SYSOUT=*
//DISK2        DD    UNIT=SYSDA,DISP=OLD,VOL=SER=222222
//TAPE2        DD    DSN=BACKUP.V222222,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN2     DD    SYSOUT=*
//DISK3        DD    UNIT=SYSDA,DISP=OLD,VOL=SER=333333
//TAPE3        DD    DSN=BACKUP.V333333.COPY1,DISP=(,CATLG),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN3     DD    SYSOUT=*
//SYSIN        DD    *
              DUMP    TYPE=DSF,DSNENQ=USE,MAXTASKS=3,MAXERR=1
              SELECT  DSN=PROD**
```

**DUMP BY
DATA SET
NAME MASK**

Dump all data sets whose names start with "SYS1." or whose names contain "SYSTEM" anywhere in the name. The three volumes will be dumped concurrently. The VTOCs will be enqueued (but not reserved) during the DUMP. A duplicate backup of the data sets from volume 'SYSRES' will be created.

```
//DUMP          EXEC  PGM=FDRDSF,REGION=3M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1        DD    UNIT=SYSDA,VOL=SER=MAST01,DISP=SHR
//TAPE1        DD    DSN=BACKUP.MAST01,UNIT=TAPE,DISP=(,KEEP)
//SYSPRIN1     DD    SYSOUT=*
//DISK2        DD    UNIT=3380,VOL=SER=SYSRES,DISP=SHR
//TAPE2        DD    DSN=BACKUP.SYSRES,UNIT=TAPE,DISP=(,KEEP)
//TAPE22       DD    DSN=BACKUP2.SYSRES,UNIT=TAPE,DISP=(,KEEP)
//SYSPRIN2     DD    SYSOUT=*
//DISKA        DD    UNIT=3380,VOL=SER=SYS002,DISP=SHR
//TAPEA        DD    DSN=BACKUP.SYS002,UNIT=TAPE,DISP=(,KEEP)
//SYSPRINA     DD    SYSOUT=*
//SYSIN        DD    *
              DUMP    TYPE=DSF,MAXTASKS=3,ENQ=ON
              SELECT  DSN=SYS1.**
              SELECT  DSN=**SYSTEM**
/*
```

20.10 CONTINUED

**A VARIETY OF
DATA SET
DUMP
COMMANDS**

The following example illustrates a variety of control statements. Data sets will be dumped from three disks to tape, one at a time. All data sets beginning with "SYS" will be dumped except those with an index level beginning with "DUMP", "PAGE", or "SWAP" in their names (since the statements are scanned in sequence, the EXCLUDE must precede its related SELECT). All data sets starting with "PAY" will have all of their allocated tracks dumped. All PO data sets with a third index of "CNTL" will be selected. Additional SELECT/EXCLUDE statements up to 500 may be specified. The volume label track (cylinder 0 track 0) will also be dumped from all disk volumes specified.

```
//DUMP          EXEC  PGM=FDRDSF,REGION=1M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    UNIT=3390,VOL=SER=SYSRES,DISP=SHR
//TAPE1         DD    UNIT=TAPE,DISP=(,CATLG),DSN=BACKUP.SYSRES
//DISKW         DD    UNIT=SYSDA,VOL=SER=MAST02,DISP=SHR
//TAPEW         DD    UNIT=TAPE,DISP=(,CATLG),DSN=BACKUP.MAST02
//DISKX         DD    UNIT=3380,VOL=SER=PAYR01,DISP=SHR
//TAPEX         DD    UNIT=TAPE,DISP=(,CATLG),DSN=BACKUP.PAYR01
//SYSIN         DD    *
                DUMP   TYPE=DSF,ENQ=RESERVE
                EXCLUDE DSN=SYS**.*DUMP**
                EXCLUDE DSN=SYS**.*PAGE**
                EXCLUDE DSN=SYS**.*SWAP**
                SELECT  DSN=SYS**
                SELECT  DSN=PAY**,*DATA=ALL
                SELECT  DSN=*.*.CNTL.*,DSORG=PO
                SELECT  FROM(CYL=0,TRK=0),TO(CYL=0,TRK=0)
/*
```

20.10 CONTINUED

**DUMP FROM
MULTIPLE
DISKS TO ONE
TAPE**

As the last step of an application system, it is desired to dump all of the data sets that have been created in this run. Tape management will expire the tape when all 3 backup GDG generations are no longer cataloged. The data sets have been allocated non-specifically across a group of three disk volumes. All data sets that start with the index level of 'BILL' will be dumped. Also a GDG relative generation will be dumped using the DD= option. Only one tape drive will be used. Only one tape volume will be used, unless the amount of dumped data requires more; the three disks will be backed up as 3 separate GDG files on the tape.

The listings from DSF will show which data set(s) were dumped from which volumes. See the example "Restore Multiple Backups" in [Section 20.12](#) for an example of restoring from a backup like this.

```
//DUMP          EXEC  PGM=FDRDSF,REGION=1M,COND=(0,NE)
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    VOL=SER=BILL01,UNIT=SYSDA,DISP=SHR
//DISK2         DD    VOL=SER=BILL02,UNIT=SYSDA,DISP=SHR
//DISK3         DD    VOL=SER=BILL03,UNIT=SYSDA,DISP=SHR
//TAPE1         DD    DSN=BILL.BACKUP1(+1),DISP=(,CATLG),
//              UNIT=TAPE,VOL=(,RETAIN),LABEL=(1,EXPDT=99000)
//TAPE2         DD    DSN=BILL.BACKUP2(+1),DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.TAPE1),LABEL=(2,EXPDT=99000)
//TAPE3         DD    DSN=BILL.BACKUP3(+1),DISP=(,CATLG),
//              VOL=REF=*.TAPE2,LABEL=(3,EXPDT=99000)
//NEWMAST       DD    DSN=ABCD.MASTER(0),DISP=SHR
//SYSIN         DD    *
                DUMP  TYPE=DSF
                SELECT DD=NEWMAST
                SELECT DSN=BILL.**
/*
```

NOTE: Since the three files on the tape must be created in ascending file number order, and the TAPEX files will be opened in the order that the DISKx DD statements appear, be sure that the LABEL=n JCL parameters specify the appropriate values. Since VOL=REF=*. will copy only the last volume serial from the previous DD statement, be sure that each TAPEX DD references the preceding one which specifies the next lower file number.

**DUMP USING
CATALOG
AND TAPEDD**

Dump several data sets. They may be on the same disk volume, or on separate disk volumes. Use the catalog to resolve the disk volume serials. Use TAPEDD= to ensure that each data set is backed up only once during this run. Use only one tape drive. Use only one tape volume, unless the volume of dumped data requires more. See the example "Restore Multiple Backups" in [Section 20.12](#) for an example of restoring from a backup like this.

```
//DUMP          EXEC  PGM=FDRDSF,REGION=2M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//DISK1         DD    DSN=USER.DATAA,DISP=SHR
//DISK2         DD    DSN=USER.DATAB,DISP=SHR
//DISK3         DD    DSN=USER.DATAC,DISP=SHR
//TAPE1         DD    DSN=BACKUP.DATAA,DISP=(,CATLG),
//              UNIT=TAPE,VOL=(,RETAIN),LABEL=1
//TAPE2         DD    DSN=BACKUP.DATAB,DISP=(,CATLG),
//              VOL=(,RETAIN,REF=*.TAPE1),LABEL=2
//TAPE3         DD    DSN=BACKUP.DATAC,DISP=(,CATLG),
//              VOL=REF=*.TAPE2,LABEL=3
//SYSIN         DD    *
                DUMP  TYPE=DSF,COMPRESS=ALL
                S     DD=DISK1,TAPEDD=1
                S     DD=DISK2,TAPEDD=2
                S     DD=DISK3,TAPEDD=3
/*
```

Note: if you want to select data sets for backup from the system catalogs without regard for which volumes they reside on, Innovation recommends that you use the ABR function of Application Backup, described in [Section 52](#).

20.11 DSF PRINT EXAMPLES

**PRINT VTOC
AND VOLUME
LABEL**

Print the VTOCs and volume labels from two disk volumes.

```
//PRINT      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//TAPE1      DD     SYSOUT=*
//TAPE2      DD     SYSOUT=*
//DISK1      DD     UNIT=DISK,VOL=SER=MYDSK1,DISP=SHR
//DISK2      DD     UNIT=DISK,VOL=SER=MYDSK2,DISP=SHR
//SYSIN      DD     *
              PRINT   TYPE=DSF,DSN=VTOC
/*
```

**PRINT
INDEXED
VTOC AND
VVDS**

Print the INDEXED VTOC (SYS1.VTOCIX.vvvvvv) and VVDS (SYS1.VVDS.vvvvvv) from a disk. SELTERR=NO is specified so that the step will end normally even if a volume contains only one or neither of those data sets.

```
//PRINT      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//TAPE1      DD     SYSOUT=*
//DISK1      DD     UNIT=DISK,VOL=SER=MYDSK1,DISP=SHR
//SYSIN      DD     *
              PRINT   TYPE=DSF,SELTERR=NO
              SELECT  DSN=SYS1.VTOCIX.**
              SELECT  DSN=SYS1.VVDS.**
/*
```

**PRINT
INDIVIDUAL
DATA SETS**

Print all data sets that have 'TEST' as the second index level. For ICF VSAM clusters, all components are printed.

```
//PRINT      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//TAPE1      DD     SYSOUT=*
//DISK1      DD     UNIT=DISK,VOL=SER=MYDISK,DISP=SHR
//SYSIN      DD     *
              PRINT   TYPE=DSF
              S        DSN=*.TEST.**
/*
```

**PRINT
ABSOLUTE
TRACK**

Print multiple absolute track segments from a single 3390 disk volume. Note that the example utilizes the 'TRK' default values for the FROM and TO operands. The output report will be written to a sequential disk file.

```
//PRINT      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//TAPE1      DD     DSN=DSF.REPORT,DISP=(,CATLG),UNIT=SYSDA,
//              SPACE=(CYL,(10,10)),DCB=BLKSIZE=3630
//DISK1      DD     UNIT=DISK,VOL=SER=MYDISK,DISP=SHR
//SYSIN      DD     *
              PRINT   TYPE=DSF
              S        FROM(CYL=0),TO(CYL=1,TRK=3)
              S        FROM(CYL=32,TRK=12),TO(CYL=35,TRK=6)
              S        FROM(CYL=21),TO(CYL=21)
/*
```

20.12 DSF NON-SMS RESTORE EXAMPLES

The following are example of DSF restores of data sets which are not SMS (System Managed Storage) managed.

**RESTORE ONE
DATA SET**

Restore one data set from a backup data set. The data set will be restored, under its original name. If the data set is currently cataloged, it will be restored to the volume to which it is cataloged; otherwise, it will be restored to the volume it was originally dumped from. In either case, if it is not in the VTOC of the disk, it will be allocated and cataloged. If it is a PS or PO data set, only its used tracks will be restored.

```
//RESTORE      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=BACKUP.VOLUMEA,DISP=SHR
//SYSIN         DD     *
                RESTORE TYPE=DSF
                SELECT   DSN=MY.TEST.DATASET
/*
```

**RESTORE
DATA SETS BY
GROUP**

Restore all the data sets which were dumped on the backup file specified by the TAPE1 DD statement, except for those data sets that start with 'SYS1.'. Active data sets will not be restored, but data sets being restored will not be enqueued. The data sets will all be restored to the volume specified by DISKx. All tracks of each data set are to be restored (the dump must also have been done with DATA=ALL). DSF will preallocate any data set not currently on the disk volume.

```
//RESTORE      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=SYSRES.BACKUP,DISP=OLD
//DISK1         DD     UNIT=3380,VOL=SER=XXXRES,DISP=SHR
//SYSIN         DD     *
                RESTORE TYPE=DSF,DATA=ALL,DSNENQ=TEST
                EXCLUDE DSN=SYS1.*
                SELECT   ALLDSN
/*
```

**RESTORE TO
NEWNAME**

Restore two data sets to new names. A data set named "LEDGER.OLDFILE" is to be renamed to "LEDGER.NEWFILE". A data set named "LEDGER.TRANSACTION" is to be restored as "LEDGER.JULY.TRANSACTION(+1)". Both will be restored to disk volume LED123. If necessary, they will be allocated and cataloged.

```
//RESTORE      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     UNIT=TAPE,DSN=BACKUP,DISP=OLD,
//              VOL=SER=12155
//SYSIN         DD     *
                RESTORE TYPE=DSF,RECAT
                SELECT   DSN=LEDGER.OLDFILE,NEWNAME=LEDGER.NEWFILE,NVOL=LED123
                SELECT   DSN=LEDGER.TRANSACTION,NEWINDEX=JULY(+1),
                        NVOL=LED123
/*
```


20.12 CONTINUED

RESTORE ICF VSAM CLUSTER Restore an existing ICF VSAM cluster from an FDR or DSF backup data set. The cluster will be restored to the volume on which it is cataloged, overlaying the existing cluster, and refreshing all associated VVDS information.

```
//RESTOREV EXEC PGM=FDRDSF,REGION=1M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPEA DD DSN=FDR.BACKUP.VSAM01,DISP=OLD
//SYSIN DD *
        RESTORE TYPE=DSF
        SELECT DSN=ICF.VSAM.CLUSTER1
/*
```

RESTORE ICF CLUSTERS TO NEWNAMES Restore two ICF VSAM clusters from an ABR-created volume backup tape. Restore "VSAM.OLD.HISTORY" to cluster "VSAM.NEW.HISTORY" which has been preallocated; DSF will do a LOCATE to get the component names of the new cluster, and restore the equivalent original components to it. Restore cluster "PAY.MASTER" to "PAY2.MASTER" which will be allocated on one of the online volume whose volser starts with "PAY"; DSF will modify all of the cluster's component names so that their first index is also "PAY2".

```
//RESTOREV EXEC PGM=FDRDSF,REGION=1M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPEA DD DSN=FDRABR.VXYZ010.C1002012,DISP=OLD,
// UNIT=TAPE,VOL=SER=121777,LABEL=3
//SYSIN DD *
        RESTORE TYPE=DSF
        SELECT DSN=VSAM.OLD.HISTORY,
              NEWNAME=VSAM.NEW.HISTORY
        SELECT DSN=PAY.MASTER,NEWI=PAY2,NVOL=PAY*
/*
```

RESTORE MULTIPLE BACKUPS Restore all the data sets that were backed up on three backup files in one execution of DSF. By default, each data set will be restored to its original disk volume, they will be allocated and cataloged if they do not currently exist. The backups were created as 3 consecutive files on the same backup tape (see "Dump from Multiple Disks to One Tape" in [Section 20.10](#) for an example). The backup data sets are cataloged, so no tape volume serials or file sequence numbers are required; since they are all on the same tape volume, only one tape drive will be allocated by MVS.

```
//RESTORE EXEC PGM=FDRDSF,REGION=1M
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DSN=BILL.BACKUP1(0),DISP=OLD,VOL=(,RETAIN)
//TAPE2 DD DSN=BILL.BACKUP2(0),DISP=OLD,VOL=(,RETAIN)
//TAPE3 DD DSN=BILL.BACKUP3(0),DISP=OLD,VOL=(,RETAIN)
//SYSIN DD *
        RESTORE TYPE=DSF
        SELECT ALLDSN
/*
```

20.12 CONTINUED

**RESTORE
MULTI-
VOLUME
DATA SETS
AND CLUSTERS**

Restore several multi-volume data sets, including non-VSAM data sets and ICF VSAM clusters. The backups pointed to by the TAPEX DDs each contain parts of these multi-volume data sets as dumped from their original volumes; the UNIT=AFF parameter insures that only one tape drive is used for all 3 backups. As DSF processes each backup, it will restore the pieces of each multi-volume data set it finds on that backup, allocating and cataloging them if necessary. Multi-volume data sets must be restored to as many different volumes as they were dumped from. If NVOL= is used, it must specify a sufficient number of volumes. A multi-volume data set will not be usable until all pieces of the data set have been restored by DSF. [See Section 80](#) for Multi-Volume VSAM Considerations; there are some restrictions.

```
//RESMULTI EXEC PGM=FDRDSF,REGION=1M
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DSN=BACKUP.VPROD01,DISP=OLD
//TAPE2 DD DSN=BACKUP.VPROD02,DISP=OLD,UNIT=AFF=TAPE1
//TAPE3 DD DSN=BACKUP.VPROD03,DISP=OLD,UNIT=AFF=TAPE1
//SYSIN DD *
        RESTORE TYPE=DSF
        SELECT DSN=NON.VSAM.MULTIO1
        SELECT DSN=NON.VSAM.MULTIO2,NVOL=(TEST01,TEST02,TEST03)
        SELECT DSN=VSAM.MULTIO3
        SELECT DSN=VSAM.MULTIO4,NVOL=TEST*
/*
```

**RESTORE TO
3390-3**

Add data sets to a 3390-3 (triple density) with DSF, in order to merge several lower density disks into one. A FDR disk-to-disk copy or a full volume restore ([See Section 10](#)) was done to move one such disk to this 3390-3. Now the contents of a 3390-1 (single density) or 3390-2 (double density) volume can be added to the 3390-3 from a FDR or DSF backup of the smaller 3390. DSF will allocate space for all restored data sets and will recatalog them to the new volume (RECAT will recatalog non-VSAM data sets, and VRECAT will scratch and uncatalog any ICF VSAM clusters before allocating them on the new volume).

```
//RESTORE EXEC PGM=FDRDSF,REGION=1M
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=3380,DISP=OLD,VOL=SER=D3390A <- - - 3390-3
//TAPE1 DD DSN=BACKUP.D3390C,DISP=OLD <- - - 3390 Backup
//SYSIN DD *
        RESTORE TYPE=DSF,RECAT,VRECAT
        EXCLUDE DSN=**CATALOG** (See Note 2, next page)
        SELECT ALLDSN (See Note 3, next page)
/*
```

20.12 CONTINUED

**RESTORE
FROM 3390-3
TO 3390-2**

Restore data sets from a backup of a 3390-3 (triple density) to three 3390 single or double density volumes. DSF will allocate as many data sets as it can on the first volume specified, and the same on the second volume, and the remainder on the third volume. The data sets will be restored to the three disks in one pass of the backup tape. DSF will allocate space for all restored data sets and will recatalog them to the new volume (RECAT will recatalog non-VSAM data sets, and VRECAT will scratch and uncatalog any ICF VSAM clusters before allocating them on the new volume). This is especially useful at a disaster/recovery site, if there are not enough high-density volumes on which to do full-volume restores.

```
//RESTORE      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//TAPE1        DD     DSN=BACKUP.D3390A,DISP=OLD
//SYSIN        DD     *
                RESTORE TYPE=DSF,RECAT,VRECAT
                SELECT  ALLDSN,NVOL=(D3390B,D3390C,D3390D)          (See Note 3)
/*
```

NOTE 2: A catalog should not be restored in the same execution as any data sets cataloged in that catalog. The catalog should be restored in a separate step.

NOTE 3: DSF will automatically exclude the SYS1.VTOCIX, SYS1.VVDS, and the ABR model DSCB from the restore, giving the warning message FDR159 REASON=3, which may be ignored. (For the VVDS, this message will be issued even if the VVDS was not selected during the backup). If you want to avoid receiving these messages and get a cleaner listing, you should specifically exclude them, e.g.,

```
EXCLUDE  DSN=SYS1.VTOCIX.*
EXCLUDE  DSN=SYS1.VVDS.*
EXCLUDE  DSN=FDRABR.V*
```

**ABSOLUTE
TRACK
RESTORE**

Restore by Absolute Track. One track will be restored from a backup on a 3480 tape drive to a 3380 disk drive, replacing that track on the output disk. A DISKx DD statement or NVOL= operand must be provided to identify the output volume.

```
//RESTORE      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=3380,DISP=OLD,VOL=SER=M33801
//TAPE1        DD     UNIT=3480,DSN=BACKUP.M33801,
//              VOL=SER=123456,DISP=OLD
//SYSIN        DD     *
                RESTORE TYPE=DSF
                SELECT  FROM(CYL=10,TRK=5),TO(CYL=10,TRK=5)
/*
```

20.13 DSF SMS RESTORE EXAMPLES

The following are examples of DSF restores on a system that has SMS (System Managed Storage) active. They illustrate the techniques for restoring SMS-managed data sets, or a combination of SMS and non-SMS data sets. An output data set will be SMS-managed if the SMS storage class ACS routine assigns a storage class to the data set or if it accepts the storage class passed to it by DSF. ACS routines are coded by each installation, so the decision on whether a data set is to be SMS-managed is a local one, and is usually out of the control of DSF (and the end-user).

RESTORE SMS-MANAGED DATA SETS

Restore a group of SMS-managed data sets from several backup data sets. The data sets will retain the SMS data class that they had when dumped. Any data sets that are currently cataloged will be restored on top of their current allocation, with their original storage and management classes. For any data sets which must be allocated, their original SMS storage and management classes will be passed as input to the SMS ACS routines, which may override them; SMS will then select a storage group and a volume. If SMS assigns a null storage group (non-SMS-managed) to any data set, that data set will be restored as non-SMS to the volume from which it was dumped.

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1        DD     DSN=BACKUP.SMS001,DISP=OLD
//TAPE2        DD     DSN=BACKUP.SMS002,DISP=OLD
//TAPE3        DD     DSN=BACKUP.SMS003,DISP=OLD
//SYSIN        DD     *
                RESTORE TYPE=DSF
                SELECT  DSN=ABC**
/*
```

RESTORE NON-SMS DATA SETS TO SMS VOLUMES

Restore a group of non-SMS-managed data sets from a FDR or DSF backup tape. Any data sets that currently exist on the volume designated by DISK1 will be restored on top of their current allocation. All other data sets will be passed to the SMS ACS storage and management class routines. If SMS assigns a storage class, SMS will then select a storage group and a volume; the data sets will not have a data class. If SMS assigns a null storage group (non-SMS-managed) to any data set, it will be allocated and cataloged as non-SMS on the DISK1 volume, which must be a non-SMS-managed volume.

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1        DD     DSN=BACKUP.PROD01,DISP=OLD
//DISK1        DD     UNIT=SYSALLDA,VOL=SER=PROD02,DISP=OLD
//SYSIN        DD     *
                RESTORE TYPE=DSF
                SELECT  DSN=*.CICS*.**
                SELECT  DSN=XYZ.**
/*
```

20.13 CONTINUED

**OVERRIDE SMS
CLASSES**

Restore a data set from a backup data set and request that the data set be SMS-managed; the original data set might be SMS-managed or not, but it is not currently cataloged. The values specified for STORCLAS= and MGMTCLAS= will be passed to the SMS ACS storage and management class routines, which may override them. If SMS assigns a storage class or accepts the class passed, SMS will then select a storage group and a volume. If SMS assigns a null storage group (non-SMS-managed) to the data set, it will be restored to the volume from which it was dumped (unless that volume was SMS-managed, since a non-SMS data set cannot be restored to a SMS volume).

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=BACKUP.PROD01,DISP=OLD
//SYSIN         DD     *
                RESTORE TYPE=DSF
                SELECT   DSN=PROD.MASTER.FILE,DATACLAS=MASTER,
                        MGMTCLAS=PERM,STORCLAS=PROD2
/*
```

**RESTORE TO
NON-SMS
VOLUME**

Restore a data set from a backup data set and request that the data set be non-SMS; the original data set might be SMS-managed or not. NULLSTORCLAS specifies that a null value will be passed to the SMS ACS storage class routine, which may override it. If SMS honors the null storage group, it will be restored to the non-SMS-managed volume TSO123; no SMS classes will be associated with it. If SMS assigns a storage class, SMS will then select a storage group and a volume.

Note: certain data set types must be SMS-managed, such as Extended Format (EF) data sets. EF data sets include striped and compressed sequential and VSAM files. Any attempt to restore these as non-SMS will fail.

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE2         DD     DSN=TEST.BACKUP.SMS123,UNIT=TAPE,
//              VOL=SER=B00123,DISP=OLD
//SYSIN         DD     *
                RESTORE TYPE=DSF,RECAT
                SELECT   DSN=USER01.ISPF.ISPPROF,NULLSTORCLAS,NVOL=TSO123
/*
```

**BYPASS ACS
ROUTINES**

An authorized user (such as a storage administrator) may need to bypass the SMS ACS routines, to force a data set to be SMS-managed or non-SMS-managed, and to specify the SMS classes to be used. In the example, TEST.DATASET1 will be unconditionally assigned the specified SMS classes; SMS will then select a storage group and a volume. TEST.DATASET2 will be forced to have a null storage class, so it will be allocated as non-SMS on the volume indicated by DISK1. See the description of the BYPASSACS operand in [Section 20.07](#) for authorization requirements.

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=BACKUP.TEST01,DISP=OLD
//DISK1         DD     UNIT=3390,VOL=SER=TEST03,DISP=OLD
//SYSIN         DD     *
                RESTORE TYPE=DSF,BYPASSACS
                SELECT   DSN=TEST.DATASET1,MGMTCLAS=TESTDS,STORCLAS=TEST
                SELECT   DSN=TEST.DATASET2,NULLSTORCLAS
/*
```

20.13 CONTINUED

**BYPASS SMS
ALLOCATION**

An authorized user (such as a storage administrator) may need to force the allocation of SMS-managed data sets onto specific volumes; normal SMS facilities do not allow you to do so, but DSF does. Each data set selected will be passed to the SMS storage and management class ACS routines; each data set to which SMS assigns a storage class will be allocated and cataloged on the volume indicated by DISKA; this volume **MUST** be an SMS-managed volume. Any data set which does not get a storage class assigned will not be restored, since only SMS data sets may be allocated on a SMS volume. See the description of the BYPASSSMS operand in [Section 20.07](#) for authorization requirements.

```
//RESTSMS      EXEC   PGM=FDRDSF,REGION=1M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//TAPEA        DD     DSN=BACKUP.DEVEL.FILES,DISP=OLD
//DISKA        DD     UNIT=DISK,VOL=SER=SMS123,DISP=OLD
//SYSIN        DD     *
                RESTORE TYPE=DSF,BYPASSSMS
                SELECT  DSN=DEVEL**
/*
```

Note: BYPASSACS and BYPASSSMS may be used together. If so, any data set which was originally SMS-managed or which has an SMS storage class specified by STORCLAS= will be directly allocated and cataloged as SMS-managed on the designated DISKx SMS-managed volume.

20.14 DSF UNLIKE DEVICE EXAMPLES

DSF can restore most types of data sets to a "unlike" disk device using a "logical" restore; for example, 3380 data sets can be restored to 3390 disks. Logical restore is automatically invoked when the output device type of a given data set differs from the type of disk backed up. Special Considerations can be found in [section 80](#) and in member UNLIKE in the FDR ICL (Installation Control Library). For ICF VSAM in particular, refer to member VSAMUNLK in the ICL for considerations on restoring clusters to smaller disk (3390 to 3380). Most of the examples in the preceding sections will work even if the output disk type is different from the type of disk backed up. Following are some specific unlike restore examples.

**RESTORE 3380
DATA SETS TO
3390**

Restore data sets from a backup of a 3380 disk to a 3390 disk. The data sets will be restored under their original names and will be cataloged to the output volume. All data sets will retain their original block sizes; they will be allocated with a size about equal in bytes to the original 3380 data set. The selected data sets may include ICF VSAM clusters and non-VSAM data sets.

```
//UNLIKE      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=BACKUP.D3380A,DISP=SHR
//SYSIN       DD     *
              RESTORE TYPE=DSF,RECAT,VRECAT
              SELECT  DSN=TEST.**,NVOL=D3390X
/*
```

**RESTORE 3390
DATA SETS TO
3380**

Restore data sets from a backup of a 3390 disk to several 3380 disks (volsers starting with "D3380"). The data sets will be restored under their original names and will be cataloged to their output volume. All PS (sequential) data sets will be reblocked to half-track blocking for maximum track utilization.

```
//UNLIKE      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=BACKUP.D3390X,DISP=SHR
//SYSIN       DD     *
              RESTORE TYPE=DSF,RECAT,BLKF=2
              SELECT  DSN=TEST.**,NVOL=D3380*
/*
```

**RESTORE
ENTIRE 3380
TO 3390**

Restore all data sets from a backup of a 3380 disk to several 3390 disks. In most cases, all the data sets on a 3380 fit on a 3390-2, but multiple 3390 output volumes are provided in case there is not room. All data sets will be restored under their original names and will be cataloged to their output volume. All PS (sequential) data sets will be reblocked to half-track blocking for maximum track utilization; PO data sets will have their block size increased to half-track for new members. The selected data sets may include ICF VSAM clusters and non-VSAM data sets. ICF VSAM clusters will be cataloged into the catalog indicated by the alias in the master catalog (by default, DSF would attempt to catalog them in their original catalog).

```
//UNLIKE      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=BACKUP.D3380X,DISP=SHR
//SYSIN       DD     *
              RESTORE TYPE=DSF,RECAT,VRECAT,BLKF=2,ICFCAT=ALIAS
              SELECT  ALLDSN,NVOL=(D3390A,D3390B,D3390C)
/*
```

20.14 CONTINUED

**RESTORE SMS-
MANAGED
DATA SET**

Restore a data set and request that the data set be SMS-managed; if SMS allocates the data set on a different disk type from its original type, a logical restore will automatically be done; if it is the same device type, a physical restore will be done.

```
//SMSUNLK EXEC PGM=FDRDSF,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPE1 DD DSN=BACKUP.PROD01,DISP=OLD
//SYSIN DD *
        RESTORE TYPE=DSF
        SELECT DSN=PROD.MASTER.FILE,STORCLAS=PROD2
/*
```


20.15 FDR INSTANTBACKUP EXAMPLES

This section contains some examples of the use of DSF with FDR InstantBackup. FDR InstantBackup is described in more detail in [Sections 25](#) through [29](#), depending on the type of disk subsystem you are using; there are more examples for each subsystem in those sections.

FDR InstantBackup is an additional cost enhancement to FDR.

INSTANT BACKUP OF VOLUME WITH SNAPSHOT

Data sets are to be backed up from an IBM RVA or StorageTek Iceberg/SVA with the Snapshot feature. The step SNAP will snap PROD01 to 01FA, creating a point-in-time image of the volume. Step BACKUP will backup the snapped copy of data sets from volume PROD01. After the backup is complete, FDR will release all of the back-end (internal disk) storage associated with the snapped volume. [See Section 26](#) for details of FDR InstantBackup for RVA/Iceberg/SVA.

```
//SNAP      EXEC      PGM=FDRSNAP
//SYSPRINT DD      SYSOUT=*
//TAPE1     DD      DUMMY
//SYSIN     DD      *
      SNAP TYPE=FDR
      MOUNT VOL=PROD01, SNAPUNIT=01FA
//BACKUP    EXEC      PGM=FDRDSF, REGION=OM, COND=(0, NE, SNAP)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1     DD      UNIT=SYSALLDA, VOL=SER=PROD01, DISP=OLD
//TAPE1     DD      DSN=BACKUP.VPROD01(+1), UNIT=TAPE,
//          DISP=(, CATLG)
//SYSIN     DD      *
      DUMP TYPE=DSF, SNAP=(USE, REL)
      SELECT DSN=PAY**
```

INSTANT BACKUP OF EMC SYMMETRIX

Data sets are to be backed up from an EMC Symmetrix with the Timefinder feature. The BCV at address 01FA has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous one-time ESTABLISH has been issued to establish that pairing. The step SPLIT will split the BCV from its online volume and wait for the split to complete, creating a point-in-time image of the volume. Step BACKUP will backup the BCV copy of data sets from volume PROD01 and re-synchronize the BCV with the online volume. [See Section 25](#) for details of FDR InstantBackup for EMC Symmetrix.

```
//SPLIT     EXEC      PGM=EMCTF
//SYSOUT    DD      SYSOUT=*
//SYSIN     DD      *
      SPLIT 1, 01FA, WAIT
//BACKUP    EXEC      PGM=FDRDSF, REGION=OM, COND=(0, NE, SPLIT)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1     DD      DSN=FDR.USE.UNIT01FA, UNIT=SYSALLDA,
//          VOL=SER=PROD01, DISP=OLD
//TAPE1     DD      DSN=BACKUP.VPROD01(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSIN     DD      *
      DUMP TYPE=DSF, BCV=(USE, RET)
      SELECT DSN=LEDGER**
```

Details of FDR InstantBackup for supported disk systems are found in [Sections 25-29](#).

20.15 CONTINUED

**DUMP WITH
HSDM
COMPRESSION**

Dump data sets from a disk volume in a disk subsystem with the HSDM (High Speed Data Mover) option installed. FDR will backup internal compressed images of each data track, reducing elapsed time up to 60%.

```
//DUMP          EXEC   PGM=FDRDSF,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1        DD     UNIT=3390,DISP=OLD,VOL=SER=ABR123
//TAPE1         DD     UNIT=TAPE,DSN=BACKUP.V123456,DISP=(,CATLG)
//SYSIN         DD     *
                DUMP    TYPE=DSF,DCT=YES
                SELECT  DSN=PROD**PAY**
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

21.01 FDRCOPY OVERVIEW

OVERVIEW The FDRCOPY program provides the ability to copy or move data sets directly from one disk to another, without an intervening backup. In a COPY operation the input data set remains on disk; in a MOVE operation the input data set is **deleted**.

In most respects FDRCOPY is similar to a DUMP with FDRDSF followed by a RESTORE, except that no backup file is involved and input and output I/O are done concurrently.

FDRCOPY can copy or move individual data sets, groups of data sets, or all of the data sets on a volume. Data sets can be selected or excluded by DSORG. Absolute tracks addresses can also be copied. Data sets can be copied or moved to the same name or to a new name. An unlimited number of data sets can be copied or moved in one execution. Data sets may be copied or moved from any number of input disk volumes, and output data sets may be directed to any number of output disk volumes. The output volume may be a different device type from the input disk volume.

FDRCOPY will allocate space for the output data set if it does not already exist on the output volume, and if requested (or by default on MOVE) will recatalog the data set to point to the new volume.

FDRCOPY can be used to convert data sets to SMS (System Managed Storage) or back to non-SMS.

FDRCOPY can also be used to reorganize (compress) PDSs in place, either by itself or in conjunction with FDRREORG, the reorganization component of FDR.

PDS REORGANIZATION FDRCOPY can reorganize PDSs (Partitioned Data Sets) in place, compressing all active members to the beginning of the PDS to allow room for new members to be added. This is similar to a IEBCOPY compress-in-place, but FDRCOPY is 50% to 90% faster than an equivalent IEBCOPY, with similar reductions in CPU and I/O resources. FDRCOPY is capable of reorganizing many PDSs on one or many volumes in one step, with just a few simple control statements. PDSs to be processed may be selected from specified volumes or selected from system catalogs, using all of the selection criteria of FDRCOPY (except that only DSORG=PO data sets will be selected).

FDRCOPY PDS REORG may be automatically invoked to process IEBCOPY compress-in-place operations without any changes to the JCL and control statements for IEBCOPY ([See Section 90](#)).

PDS reorganization is requested by the REORG statement, as described in [Section 32.01](#).

However, the REORG function is enabled only if your installation is licensed for FDRREORG, a separately-priced component of FDR.

FDRCOPY also supports simulation of PDS reorganization, via the SIMREORG control statement. SIMREORG must still read the PDSs but will only simulate rewriting the data, reporting on the number of tracks that will be reclaimed by a real REORG.

FDRREORG ([see Section 30](#)) reorganizes VSAM data sets, IAM (Innovation Access Method) data sets, and PDSs, using FDRCOPY internally for the PDS reorganization. If you execute PDS reorganization through FDRREORG, it enhances the reorganization process by providing additional selection criteria and retrying unavailable data sets, among others. However, these enhancements may add additional overhead; executing FDRCOPY REORG directly will provide the fastest PDS reorganization.

**SPECIFYING
DATA SETS TO
BE COPIED OR
MOVED**

FDRCOPY copies or moves user-specified data sets from one or more input disks to one or more output disks. The data sets to be processed can be specified by individual name, or groups of data sets may be selected using a generic data set name filter, or the ALLDSN operand may be used to copy or move all of the data sets on a volume. The data sets may be selected by scanning the VTOCs of volumes specified by the user, or data sets may be selected from system catalogs. Data sets can be selected or excluded by DSORG.

By default data sets are copied or moved under their original name but you can cause the output data sets to get a new name. COPY will usually use a new name (so that both the original and copied data set can be cataloged) while MOVE usually uses the original name. For a COPY of a VSAM cluster, the output data set must either have a different name or be cataloged in a different catalog; on a VSAM MOVE the output data set usually has the original name.

21.01 CONTINUED . . .

**DATA SET
COPY/MOVE**

During a COPY or MOVE:

- the data tracks of the selected data sets may be copied to a different physical location (cylinder and head address) than the original tracks occupied. In other words, you do not have to be concerned about the location of the output data set or its extent count.
- data sets and clusters may be renamed during the COPY/MOVE.
- a target output volume will be selected for each data set being copied. If you are copying to a new name, this is the input volume by default, but you can override the target. Target volume selection rules are detailed below.
- if an output data set (either the original data set name or the new name) already exists on its target volume, and is large enough, FDRCOPY will replace the contents of the existing data set and update its describing information (VTOC and VVDS information). However, if you specify the PRESTAGE operand, FDRCOPY will bypass the copy of a preexisting data set.
- if an output data set does not exist on its target volume, FDRCOPY will allocate it, automatically making it large enough to hold the contents of the input data set. Usually FDRCOPY will also catalog the output data set at the end of the copy (if the output data set is already cataloged to other than the input volume, FDRCOPY will not catalog it unless you instruct it to do so). VSAM clusters must be cataloged when they are allocated, so a VSAM allocation will fail if it is already cataloged.
- FDRCOPY usually copies data sets to the same device type they were backed up from (e.g., from 3390 to 3390). This is called a like device or physical copy. The size (number of cylinders) on the original disk volume and the output volume is not important, as long as the output has enough free space to hold the output data sets being allocated. In a like copy, the original data tracks of the selected data sets are copied exactly as they were backed up (but there is an option to reblock certain data set types). FDRCOPY will allocate an output data set of the same size as each input data set, unless you override the size with various operands.
- FDRCOPY can also copy data sets to a different device type (such as 3380 to 3390), with some restrictions. This is called an unlike device or logical copy since the data records of the original data sets are usually reformatted to make better use of the track capacity of the output disk. On an unlike device copy, the input data set may require a different amount of disk space on the output disk, depending on data set organization, block sizes and reblocking; FDRCOPY will calculate an appropriate amount of disk space for each data set.
- Output data sets may be directed to various output disk volumes. Data sets from one input volume can be copied to many output disks concurrently. The input and output volumes can be identified by JCL or by FDRCOPY control statements. You can also identify a list or group of volumes as the target volume; FDRCOPY will find one with sufficient space for the data set. However, a single volume data set must be copied to a single volume.
- Multi-volume data sets can be copied, but only to the same number of volumes they currently occupy when dumped. Multi-volume VSAM is handled, but only when copied to the same device type with a physical copy. One exception: if a cluster is multi-volume because it has a data component on one volume and an index component on another volume, both components can be copied to one volume.
- At the end of the copy from each input volume, FDRCOPY will update the DSCB of each output data set and, for VSAM and SMS-managed data sets, its VVDS entry, so that they properly describe the data that was copied. If a copy is interrupted or cancelled, this update will not be done and the data sets will probably be unusable even though all data tracks were copied.
- If FDRCOPY allocated an output data set and the copy of that data set gets errors, such as disk I/O errors, the data set will be deleted from disk, to avoid leaving unusable and uncataloged data sets.

21.01 CONTINUED . . .

**TARGET
VOLUME
SELECTION**

While processing the data sets selected from a given input disk volume, FDRCOPY can copy/move them to one or more disk volumes concurrently. The target disk volume will be selected for each data set by the following rules:

- If the NVOL= operand was specified on the SELECT statement which selected this data set, that volume or volumes will be used. See the description of NVOL= in [Section 21.05](#) for details of target volume selection.
- If the input volume being processed was specified by a DISKx DD statement in the FDRCOPY JCL, and there is a matching TAPEx DD pointing to a disk volume, then the disk to which it points will be selected for output.
- If you are copying or moving a data set to a new name, and the output data set name is cataloged, then the volume to which it is cataloged will be chosen. If the data set is cataloged as being on multiple volume serials, then the volser will be selected from that list based on the volume sequence number in the F1 DSCB (field DS1VOLSQ) of the input data set. When copying or moving a data set under its original name, FDRCOPY will not check if the data set is cataloged to another volume, in order to prevent accidentally overlaying a current cataloged data set with a copy of an uncataloged data set with the same name.
- If none of the above apply, then the data set will be copied or moved to the input volume. However, any attempt to copy or move a data set back on top of itself will be rejected.

If SMS (System Managed Storage) is active on this system, and the data set does not already exist on the volume selected by the rules above, SMS is invoked to decide if the data set should be SMS-managed. If so, SMS will select an output volume. SMS rules are detailed later in this section.

**CATALOGING
NON-VSAM
OUTPUT DATA
SETS**

When copying or moving single-volume non-VSAM data sets:

- On a COPY, if FDRCOPY allocates space for the output data set, then the default is that FDRCOPY will catalog the data set to the output volume, unless the data set was already cataloged. (For a copy to the original name, if the input data set is cataloged, the output data set will not be cataloged).
- On a MOVE the default is that if FDRCOPY allocates space for the output data set, then FDRCOPY will catalog the data set to the output volume unless the data set was cataloged to a volume other than the input volume before the operation. (For a move to the original name, if the input data set is cataloged to the input volume, it will be recataloged to the output volume).
- If the RECAT operand is specified, it will always catalog the newly allocated data set to the output volume, whether the data set was cataloged to the input volume before the operation, or cataloged elsewhere, or not cataloged at all. This allows you to insure that the catalog will point to the copied data set.
- If the NOCAT operand is specified then FDRCOPY will not catalog the newly allocated data set. This is ignored for SMS-managed data sets which must always be cataloged.
- Cataloging or recataloging of non-VSAM data sets occurs at the end of the copy/move, and is bypassed if any errors have occurred for a given data set.
- If the output data set exists on disk before the restore, then FDRCOPY will not update the catalog, unless the CATIFALLOP operand is specified. CATIFALLOP causes FDRCOPY to apply the same rules described above for data sets allocated by FDRCOPY.

21.01 CONTINUED . . .

**CATALOGING
NON-VSAM
OUTPUT DATA
SETS
(continued)**

- STEPCAT DDs are supported; if present, only the STEPCAT catalog (or the first catalog in the STEPCAT concatenation) will be searched and updated for non-VSAM data sets. The target catalog for VSAM clusters is controlled by the ICFCAT= operand described in [Section 20.08](#). STEPCAT is not supported if SMS-managed data sets must be allocated.

NOTE: If a COPY/MOVE is interrupted by an ABEND or system crash, then the output data sets will be left on disk, but will not be cataloged. You should resubmit the FDRCOPY job, to insure that all data tracks are copied, but the output data sets will already be allocated so FDRCOPY will not catalog them unless you specify the operand CATIFALLOC.

For multi-volume non-VSAM data sets, the above rules for single-volume data sets apply, with the following modifications:

If the data set is not already cataloged (e.g., MOVE or COPY to a new name), FDRCOPY will create a new multi-volume catalog entry with the current output volume in the proper slot as indicated by the volume sequence number in the DSCB. If the volume sequence number is higher than 1, FDRCOPY will fill in the preceding slots in the catalog entry with a dummy volume serial of "####nn".

If the data set is already cataloged FDRCOPY will update the catalog entry by putting the current output volume into the proper slot as indicated by the volume sequence number in the DSCB. On a COPY without RECAT, FDRCOPY will update the catalog entry only if the slot for this output volume contains the dummy volume serial of "####nn"; otherwise a warning message will be issued.

So, when doing a MOVE, or a COPY with RECAT, of a cataloged multi-volume data set to the same name, the resulting data set and catalog entry will be usable and correct, whether the data set is moved from one, some, or all of its original volumes. When doing a MOVE or COPY of a cataloged multi-volume data set to a new name, it will be necessary to move or copy the data set from all of its volumes to get a usable data set and catalog entry. If a multi-volume data set is being moved/ copied to a new device type, it will not be usable until all pieces are on the same device type. If you do not copy one or more pieces of the data set, the catalog entry will be inaccurate and the data set will not be usable (it may contain ####nn or original volsers). It is your responsibility to insure that the data set is copied from all the volumes it occupies; luckily, this is easy with the CATDSN= operand, selecting data sets from the catalog.

All of these rules sound complicated, but they are designed to automatically do the right things when cataloging non-VSAM data sets. You need to be concerned only if a data set being copied may already be cataloged to other than the input volume; if so, specify RECAT if you want the newly copied data set to be the cataloged data set. Specify CATIFALLOC if there is a chance that the output data set is already allocated but is not currently cataloged (and you want it to be the cataloged version).

SMS-managed data sets are always cataloged, so a copy/move to an SMS data set requiring allocation will fail if the data set cannot be cataloged. When moving SMS-managed data sets they are automatically recataloged to the new volume, but for a copy to a new name that already exists somewhere, RECAT may be required. NOCAT is ignored for SMS data sets.

21.01 CONTINUED . . .

**DELETING THE
INPUT DATA
SET**

At the end of a MOVE operation, the input data sets are deleted. Deletion is bypassed for a given data set if any errors have occurred for that data set.

For a multi-volume data set, only the portion on the current input volume will be deleted. Portions on other volumes will remain on disk, unless they are also MOVED.

FDRCOPY will not delete multi-volume VSAM data sets. It will also not delete VSAM clusters which are not properly cataloged or not cataloged at all.

On a COPY operation, the input data set remains on disk (although the catalog is updated to point to the output data set, if the RECAT operand is specified and the output data set has the same name).

**ABSOLUTE
TRACK
ADDRESS
OPTION**

If the user specifies physical track starting and ending addresses, FDRCOPY will COPY the requested physical segment from the input volume to the same physical location on the output volume. This operation will not update the VTOC on the output volume to reflect the data sets that have been affected.

A MOVE operation for absolute track addresses will be treated as a COPY. Since the operation is not associated with a data set, there is no input data set to delete.

Extreme caution should be used when copying tracks within the VTOC, VTOC index, or VVDS, or the volume label track, since an error may make data sets or the entire pack unusable.

**ICF VSAM
SUPPORT**

FDRCOPY supports ICF VSAM files using the base cluster name. FDRCOPY will copy or move all of the components associated with the cluster on a volume, including alternate indexes and key ranges. In addition FDRCOPY will copy or move the VVR information found in the VVDS. An ICF VSAM cluster can be selected by individual name or as part of a data set group, or will be included when the ALLDSN operand is used to copy or move all of the data sets on a volume. If the output cluster does not exist on disk before the copy or move, then FDRCOPY will allocate space for it by invoking SVC 26 to do a DEFINE.

When FDRCOPY does a DEFINE for VSAM, some of the information that exists only in the catalog will be copied from the input cluster, including candidate volumes, expiration date, owner ID, and passwords and other VSAM security fields. The creation date will be copied on a MOVE operation; on a COPY operation the creation date of the output cluster will be the run date. If the input cluster is protected by a discrete RACF profile, then the output cluster will also be protected by a discrete RACF profile, which will be created by using the profile of the input cluster as a MODEL. If the input profile name does not actually exist, the output cluster will be created with no discrete profile. Path names for alternate indexes **will** be copied or moved.

For VSAM, all data sets must be cataloged; the NOCAT, RECAT, and CATIFALLOC operands are ignored.

FDRCOPY copies track by track; the cluster is not reorganized.

[See Section 80.13](#), VSAM Special Considerations, for further details.

21.01 CONTINUED . . .

ICF VSAM – MOVE A MOVE operation for a VSAM cluster can use either the same name or a new name. On a MOVE to the original name, FDRCOPY will temporarily DEFINE the output cluster with an inserted index level of “.Txxxx”, where xxxx is a timestamp consisting of four hexadecimal digits. “.Txxxx” will be inserted as the next-to-last index level of the temporary name for the output cluster and for each of its components. If the original name is longer than 38 characters, characters will be removed from the original name as necessary, preceding the point of insertion. If these temporary names would violate your installation naming standards, please contact INNOVATION for assistance.

At the end of a MOVE, if no errors have been encountered, FDRCOPY will DELETE the input cluster. If the MOVE was to the original name, FDRCOPY will then ALTER the names of the output cluster and its components back to the original names.

If errors are encountered during a MOVE, before the input cluster is deleted, FDRCOPY will attempt to DELETE the output cluster; if it is successful, no action is required and the input cluster is unchanged. If it fails, then both the input cluster and the output cluster will remain on disk and you will need to manually clean up:

- If the MOVE was to the original name, then the output cluster will have the temporary name. If the move is re-attempted, a new output cluster will be created with a different timestamp in its temporary name. After the earlier output cluster is no longer needed for investigation of the error, you should DELETE it with IDCAMS.
- If the MOVE was to a new name, then the output cluster will have the new name. If the MOVE is re-attempted, FDRCOPY will find and reuse the earlier output cluster.
- If errors are encountered during a MOVE, while the input cluster is being deleted, then the input cluster may or may not still be on disk, as indicated by the error codes, and the output cluster will be on disk. If the MOVE was to the original name, then the output cluster will have the temporary name. You can use IDCAMS DELETE and ALTER to clean up this condition.

If errors are encountered during a MOVE to the original name, while the output cluster is being renamed, then the input cluster will be gone, and the output cluster will remain on disk; some parts of the output cluster may have been renamed to the original name. You can use IDCAMS ALTER to clean up this condition.

A MOVE of a multi-volume ICF VSAM cluster is not supported. A MOVE of a KSDS with keyranges, even though they are on the same volume as the base cluster, can only be done to a new name. Multi-volume VSAM can be moved by dumping all components with DSF, deleting the cluster, and restoring to new volumes ([see Section 20](#)).

ICF VSAM – COPY A COPY operation for a VSAM cluster must always specify either a new name for the output cluster and all its components, or a different catalog (ICFCAT=). The reason is that all ICF VSAM files are cataloged as soon as they are created, and there cannot be two catalog entries for the same cluster name in the same catalog.

A COPY of a multi-volume ICF VSAM cluster is supported in most cases. [See Section 80](#) for Multi-Volume VSAM Considerations.

ICF VSAM – NEW NAMES For a VSAM COPY/MOVE to a new name, it is recommended that you use the NEWGROUP= or NEWINDEX= operands to rename the cluster; the changes specified will be made to the output cluster name, all its component names, and any alternate index (AIX) cluster names. For a cluster with no AIXs, you can specify NEWNAME=; the cluster will be renamed and IBM defaults will be used to name the components (which usually involves adding “.DATA” or “.INDEX” to the cluster name).

ICF VSAM catalogs and VVDSs cannot be copied to a NEWNAME.

ICF CATALOG FDRCOPY cannot copy or move an ICF VSAM catalog (BCS), because an ICF catalog cannot be copied or moved to a new name, even temporarily. You can move an ICF catalog by dumping it with DSF, deleting the old catalog with the IDCAMS command “DELETE catname RECOVERY”, and restoring the catalog to the new volume ([see Section 20](#)).

VVDS FDRCOPY cannot copy or move a VVDS (VSAM Volume Data Set, SYS1.VVDS.Vserial) as an entity. Instead, FDRCOPY updates the VVDS on the output volume to reflect the attributes of the clusters that have been copied or moved.

21.01 CONTINUED . . .

SMS SUPPORT On a system with IBM's SMS (System Managed Storage) active, FDRCOPY supports copying and moving SMS-managed data sets, and conversion of non-SMS data sets to SMS management and back again.

When processing input data sets which are SMS-managed, FDRCOPY will extract SMS information from the VVDS and VTOC for both VSAM and non-VSAM data sets. This includes SMS class information (storage, management, and data classes), and SMS indicators.

When output data sets must be allocated on a SMS system, SMS will be invoked for every such data set, to decide if the data set should be managed by SMS, or allocated as non-SMS. The SMS storage class and management class ACS (Automatic Class Selection) routines will be invoked; they will be passed input class names:

- if the user specified STORCLAS=, NULLSTORCLAS, MGMTCLAS=, or NULLMGMTCLAS, those overriding values will be used.
- if the storage class or management class (or both) were not overridden by the user, the class associated with the input data set will be used (if the input data set was not SMS-managed, a null class will be passed).

The ACS routines may accept those classes, or override them with different values or even null values.

If SMS assigns a storage class to a data set, it will be SMS-managed; SMS will be invoked again to allocate the data set on a volume chosen by SMS. If no storage class is assigned, FDRCOPY will allocate the data set on a non-SMS volume (a target non-SMS volume must be indicated by a TAPEX DD statement or a NVOL= operand on the SELECT statement). Even if data sets are allocated by SMS on a number of different volumes, FDRCOPY will output those data sets in one pass of the input volume.

So, FDRCOPY can be used to convert data sets to SMS management, simply by updating the SMS ACS routines to assign storage classes to the output data sets, or by specifying a storage class via the STORCLAS= operand on the SELECT statement. Data sets can be converted back to non-SMS if the ACS routines assign no storage class or if the NULLSTORCLAS operand is specified. FDRCOPY can also be used to move SMS-managed data sets to new volumes if the SMS storage groups are redefined.

Storage administrators, with proper authority, can override or bypass many of the SMS functions, to directly specify SMS classes, or to specify the volume serial on which SMS data sets are to be placed, by use of the BYPASSACS and BYPASSSMS operands on the COPY/MOVE statement.

More detail on SMS support is in [Section 70](#).

**PROCESING
SPECIFIC
TYPES OF
DATA SETS**

Details of processing for various types of data sets are in [Section 80](#).

21.01 CONTINUED . . .

**FDR
INSTANTBACKUP**

If you are also licensed for FDR InstantBackup, it enhances FDRCOPY to automatically use the hardware features of certain disk subsystems to move data tracks from one volume to another without the channel overhead of sending the data to and from the CPU. This currently works on IBM RVA and StorageTek Iceberg/SVA systems with the Snapshot feature and EMC Symmetrix systems with the Timefinder feature. When FDRCOPY detects that the input and output disks for any data set are in the same disk subsystem, and that subsystem is enabled one of the supported hardware features, it will be automatically be used. The elapsed time of the COPY or MOVE will be significantly reduced. FDR InstantBackup for Symmetrix and RVA/Iceberg/SVA is described in [Sections 25](#) and [26](#).

FDR InstantBackup also enhances FDRCOPY to provide the ability to take copy data sets which are frozen at a given point-in-time. This works with almost any disk subsystem. It allows you to capture a point-in-time image of an online disk to an offline disk, effectively preserving the image of the online disk at that point-in-time. FDR InstantBackup can then read the offline disk and create the required backups or copies, without relabeling the disk or bringing it online. FDR InstantBackup for various hardware platforms is described in [Sections 25-29](#).

21.02 FDRCOPY PROCESSING OPTIONS AND REQUIREMENTS

**FDRCOPY
OPERATIONS**

A **COPY** operation makes copies of the selected input data sets. In most cases, the output data sets will be given new names, to avoid creating duplicate data sets names (since only one copy of a data set name can be cataloged). You can use COPY to create duplicate names for non-VSAM data sets, such as when cloning a MVS residence volume, but the output data sets will not be cataloged unless you force them to catalog (making the original data sets uncataloged). For ICF VSAM, the output data set must have a name different from the input unless you catalog it in a different ICF catalog.

A **MOVE** operation is similar to a COPY, except that the input data sets are deleted when successfully copied. MOVE allows you to move selected data sets to new volumes. In most cases, the output data sets will retain their original names and the catalog will be updated to point to the new location (if the catalog pointed to the input data set). For ICF VSAM, a temporary name will be used for the cluster and its components; it will be renamed after the input data set is deleted.

A **REORG** operation reorganizes a PDS (Partitioned Data Set) in place. This is similar to a IEBCOPY COMPRESS but is much faster. REORG is functional only if you are licensed for FDRREORG, a separately-priced component of FDR ([See Section 30](#)).

**MEMORY
REQUIREMENTS**

Most FDRCOPY jobs will run in a region of 2000K of below-the-16MB-line storage. If a very large number of data sets or VSAM clusters are to be processed, more storage may be required. Some logical copy operations may also require additional memory. FDRCOPY memory requirements are fixed for a given function; if too large a region is provided FDRCOPY will not make use of the excess memory, but it will fail if too small a region is given. So, we recommend that you specify as large a region as possible (REGION=0M will allocate the largest possible region).

On ESA and OS/390 systems, some FDRCOPY functions use memory above the 16MB line. The normal default region of 32M above the line is usually adequate, but if your default is smaller or more is required, specify a larger value (e.g., REGION=64M allocates 64MB above the line and the largest possible region below).

**ERROR
DETECTION**

If any of the selected data sets are not found on any input disk or have errors (such as I/O errors), FDRCOPY will continue the copy/move operation for the remainder of the data sets. An error message will identify the failing data set and a U0888 abend will be issued for the step at completion to call attention to the error.

21.02 CONTINUED . . .

SECURITY Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".
WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system which supports the SAF (Security Authorization Facility) interface such as ACF2 or TOP SECRET, you can enable the ALLCALL option. For FDRCOPY this results in these security checks:

- for input data sets for COPY and MOVE, FDRCOPY will always check to see if your userid has at least READ authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority, then FDRCOPY will check if you have at least READ authority under the DATASET security class to every data set being copied. Any data sets to which you are not authorized will be bypassed with an error message. However, on a MOVE, the delete of the input data set will fail if you do not have ALTER authority to the data set.
- for output data sets for COPY or MOVE, when data sets must be allocated on the output volume, the system services called by FDRCOPY to allocate the data set will do their own checks to be sure you are authorized to do so; this is always done regardless of FDR security options. For pre-existing output data sets, FDRCOPY will check to see if your userid has UPDATE authority to the entire output volume. If you do have this volume authority, no additional checks are done on that output volume. If you do not have volume authority, then FDRCOPY will check if you have UPDATE authority under the DATASET security class to every pre-allocated output data set. Any data sets to which you are not authorized will be bypassed with an error message.
- if the COPY or MOVE includes absolute track address selections (FROM/TO), only the DASDVOL checks described above are done for those absolute track ranges. If volume security rules are not defined, any user can do absolute track operations on the unprotected volumes, so Innovation recommends always defining such rules. You can disable absolute track operations by an option in the FDR Global Option Table ([Section 90](#)).

CONTINUED . . .

21.02 CONTINUED . . .

**DATA SET
ENQUEUE
OPTION**

You can request, via the DSNENQ= operand, that each data set being copied or moved be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name of SYSDSN and a minor name of the data set name itself; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. This test is done for selected input data sets selected by FDRCOPY and also for output data sets if you are assigning them a new name. Note that FDR cannot tell if an active data set is being used for input or output by the other task. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations. If you have requested or defaulted to data set ENQs, any data set that is in use will cause a FDR158 warning message to be printed; this will set the job error flag and will cause a U0888 abend when the step is complete (see "Step Termination" below). If you don't want in-use data sets to be considered an error, specify the ENQERR=NO operand; this prints the FDR158 message without setting the error flag.

You can request that inactive data sets be enqueued to FDRCOPY during the copy/move, to insure that no other job or TSO user can access the data set until the copy is done. This is the default.

In-use input data sets will not be copied by default. If you wish to copy active data sets, specify the ENQERR=PROCESS operand, but you must be aware that the copies of data sets which are being updated during the copy may be unusable, depending on the nature and format of the data.

Output data sets will be enqueued only if they are a new name, different from the input. FDRCOPY will attempt to enqueue any new name data sets that it allocates on the output disks, to insure that no other task tries to use them until the copy is complete, but if the ENQ fails, the data set is still copied. But for existing pre-allocated new name data sets, if the ENQ fails, the copy will be bypassed.

The DSNENQ= operand has 4 possible values:

USE – data sets will be enqueued for the duration of the copy from an input disk volume. For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option and it is the default for FDRCOPY.

TEST – data sets will only be tested to see if they are enqueued to another task at the time that the copy from this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the copy is proceeding.

HAVE – The data sets will be enqueued for the duration of the copy. If a data set is in use, the MVS operator must interact with FDR to decide how to proceed; a message (FDRW27) is issued to the MVS console, and the operator can respond:

WAIT - wait for the data set to become available; if it is not eventually dequeued, the FDR job may time out, so the operator must not reply WAIT for data sets in use by long-running jobs or tasks such as transaction processing systems like CICS.

NOWAIT - do not enqueue the data set. The FDR158 warning message is issued.

RETRY - try the enqueue again. If it fails again, the FDRW27 message is reissued.

NONE – No data set ENQ will be issued. This can be used for data sets which you know will appear to be active, such as data sets on a MVS residence volume.

NOTE: If a data set name appears in a DD statement with DISP=SHR within the FDR job (not necessarily in the FDR step), FDR will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the FDR job ends.

CAUTION: *This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.*

21.02 CONTINUED . . .

**DATA SET
ENQUEUE
OPTION
(continued)**

The default of DSNENQ=USE prevents other tasks from updating (or reading) data sets being copied or moved, while bypassing data sets which are in use. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

Note: a MOVE will delete the input data sets using IBM SCRATCH or DELETE services. These services may do their own SYSDSN ENQ on the data sets and may fail the deletion if the data set is in use.

If HFS=QUIESCE is specified, special processing is done for HFS data sets (Hierarchical File System, used by OS/390 Unix System Services). If the SYSDSN ENQ cannot be acquired, this may mean that the file system is mounted to OS/390 Unix, so FDR will attempt to quiesce the file system during the backup. Details on the quiesce function are found in [Section 80.11](#).

**STEP
TERMINATION**

If no errors occur during the execution of FDRCOPY, the FDRCOPY jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

- Some errors are critical. The jobstep ends immediately with a user ABEND.
- Some errors are critical only to a particular operation. For example, during a copy, some errors cause the copy from a particular disk to terminate immediately, but FDRCOPY may continue and attempt to copy other disks requested in the same step.
- Some errors are non-critical and the messages are warnings only. FDRCOPY will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with a U0888 abend to call your attention to the errors. Remember that a U0888 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer not to get a U0888 abend on a non-terminating error, the FDRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice ([see Section 90](#)).

21.03 FDRCOPY JOB CONTROL REQUIREMENTS

To execute FDRCOPY, the following JCL statements are required:

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRCOPY resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRCOPY), region requirement (REGION=, [see Section 21.02](#)), and optional PARM= operand.

If a PARM field is specified, FDRCOPY will use data specified as the first control statement, which must be a COPY or MOVE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRCOPY,PARM='MOVE TYPE=DSF,DSNENQ=HAVE'
//FDR EXEC PGM=FDRCOPY,PARM='COPY TYPE=DSF/ SELECT DSN=A.B.C,NEWN=X.Y.Z'
```

If FDRCOPY is invoked from another program, you can pass control statements using IBM's convention for passing data from the PARM field; contact Innovation for assistance if you wish to do this.

STEPCAT or JOBCAT DD STATEMENT For non-VSAM data sets, when FDRCOPY must catalog non-VSAM data sets, they will be cataloged in that user catalog instead of in the system catalog with the matching alias. This might be useful when you are creating a test system, to catalog test copies of production data sets in a test catalog.

However, there is one important exception: if the data set being cataloged is a GDG generation, the STEPCAT/JOBCAT must contain a GDG base for that GDG; if not, it will ignore the STEPCAT/JOBCAT and catalog into the regular aliased catalog, possibly deleting other valid generations.

For VSAM clusters, the target catalog is controlled by the ICFCAT= operand, described in [Section 21.04](#). With the proper ICFCAT= option, the STEPCAT/JOBCAT may be honored.

STEPCAT/JOBCAT should not be used if any data sets being copied are SMS-managed.

SYSPRINT DD STATEMENT Specifies the output message data set. It must be present and is usually a SYSOUT data set but it may be assigned to disk or tape. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape. SYSPRINT will receive messages from the restore (output) side of the copy/move/reorg, and will also receive messages from the dump (input) side if SYSPRINx DD statements are not provided.

SYSPRINx DD STATEMENT Specifies an optional output message data set (see SYSPRINT for details). If optional DISKx DD statements were used to specify input disks, you may optionally provide a SYSPRINx for each DISKx DD in the step. SYSPRINx will receive messages from the dump (input) side of FDRCOPY, but these messages will go to SYSPRINT if SYSPRINx is omitted (in this case, some non-essential dump messages will be omitted).

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```


21.03 CONTINUED . . .

DISKx DD STATEMENT Optionally specifies an input disk volume, or, for REORG, a volume containing PDSs to be reorganized. The format will be:

```
//DISK1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial assigned of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). Only a single disk volume serial may be specified. You may use either DISP=OLD or DISP=SHR; it makes no difference.

"x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character and must have a corresponding TAPEx statement, so there can be a maximum of 39 DISKx DDs. If DUMMY is specified, this DD statement will be ignored.

FDRCOPY will process all disk volumes specified by DISKx statements in the FDRCOPY step in the order that the DDs appear in the JCL; however, any disk volumes identified by VOL= operands on SELECT statements, or selected from the catalog by CATDSN= will be automatically dynamically allocated by FDRCOPY if there is no DISKx DD statement for them, **so DISKx DDs are not required.**

For FDR InstantBackup

If you are also licensed for FDR InstantBackup, you can direct FDRCOPY to copy from an offline point-in-time image of the volume to be backed up. This allows you to capture that point-in-time image and copy data sets from it at your leisure, even while updates are being done to the live, online disk. This may require special options on a DISKx DD statement. FDR InstantBackup is described in [Sections 25](#) through [29](#), with separate sections for each hardware platform on which FDR InstantBackup is supported.

TAPEx DD STATEMENT Optionally specifies an output disk volume, by the same rules as DISKx above, e.g.,

```
//TAPE1 DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

TAPEx will be ignored if there is no matching DISKx DD statement, or if it specifies DUMMY.

Although TAPEx DDs are optional, if present they specify a default target output volume for data sets moved or copied from the disk volume specified by DISKx. That default may be overridden by the NVOL= operand on the SELECT statement which selected the data set, or by SMS.

SYSIN DD STATEMENT Specifies a data set containing the control statements for FDRCOPY. Usually a DD * data set. It is required, but if control statements were provided in the EXEC PARM=, it can be DUMMY.

21.04 FDRCOPY COPY/MOVE STATEMENT

COPY	TYPE=DSF	,ICFCORE=nnnnnn
MOVE	,BLKF=nn	,MAXCARDS=nnnnnn
	,BYPASSACS	,MAXERR=nnnn
	,BYPASSSMS	,NOCAT
	,CATIFALLOC	,RECAT
	,DATA=ALL USED	,PRESTAGE
	,DSNENQ=NONE TEST USE HAVE	,RLSE
	,ENQERR=NO	,%FREE=nn
	,ENQERR=BYPASS PROCESS	,SELTERR=NO YES
	,HFS=QUIESCE	,MSGDG=DEFERRED ACTIVE ROLLEDOFF INPUT
	,ICFCAT=ORIGINAL STEP CAT ALIAS	,SNAP=(USE,REL)

COPY STATEMENT The COPY statement specifies that FDRCOPY is to perform a copy operation, creating a duplicate of the selected input data sets. The original data sets will not be modified in any way. However, if the RECAT option is specified and the output data set is not being given a newname, the catalog will be updated to point to the data set on the output volume serial, leaving the input data set uncataloged. ICF VSAM data sets can be copied only if a new output data set name is specified.

MOVE STATEMENT The MOVE statement specifies that FDRCOPY is to perform a move operation, moving the selected input data sets to a new disk volume. For each successfully moved data set, the input data set will be deleted and uncataloged, and the catalog will be updated to point to the output data set. Moves usually use the same data set name for output, making them a true move. But you can specify a new output data set name; at the end of the move the input data set is still deleted and only the renamed output data set is left.

OPERANDS

TYPE=DSF Specifies that a data set copy/move is to be performed. It is required.

BLKF= Changes the blocksize (BLKSIZE) of certain fixed (RECFM=FB) and variable (RECFM=VB) data sets. BLKF= specifies the number of blocks per track (1 to 10) and is used to calculate a new blocksize based on the output track size (BLKF=1 results in BLKSIZE=32760). For RECFM=FB data sets the blocksize will be rounded down to a multiple of the logical record length (LRECL).

DSORG=PS (sequential) data sets will be physically reblocked during the copy; the new blocksize must be larger than the original blocksize except when restoring from 3390 to 3380.

DSORG=PO (partitioned) data sets will not be reblocked but the new blocksize will be used for new members; it must always be larger than the original blocksize.

Other data set types, including UN (undefined), DA (direct access), PSE (enhanced format sequential), POE (PDSE), HFS and all types of VSAM, will not be reblocked.

Default: data sets are not reblocked during copy; all original blocks will be copied without change, although they may be written to new locations. BLKF= is usually used when copy to an unlike device type (e.g., 3380 to 3390) to improve utilization of the new track size but can also be used during like device copies.

21.04 CONTINUED . . .

BYPASSACS On a system with SMS (System Managed Storage) active, the SMS ACS (Automatic Class Selection) routines are not to be invoked for data sets which must be allocated. If a data set has a SMS storage class assigned (see STORCLAS= in [Section 20.09](#)) it will be SMS-managed, and SMS will be invoked to allocate the data set on an SMS-chosen volume, but SMS will not be allowed to override the storage class or management class assigned to the data set.

Default: on an SMS system, the SMS ACS routines will be invoked for every data set which has to be allocated. The assigned storage and management classes will be passed to those routines, which can approve or override them. A data set will be passed to SMS for allocation if the storage class ACS routine assigns a storage class to the data set.

BYPASSSMS On a system with SMS (System Managed Storage) active, SMS data sets will be directly allocated on SMS-managed volumes, bypassing normal SMS storage group and volume selection. The selected output volume must be a SMS-managed disk volume, and the data sets being copied must have a SMS storage class assigned (see BYPASSACS above and STORCLAS= in [Section 20.09](#)). The data sets will be allocated and cataloged according to SMS standards.

Normal SMS facilities do not allow allocation of data sets on specific volume serials, but BYPASSSMS will do so, allowing data sets to be located for performance or other reasons. Note that if BYPASSACS is also specified, the assigned SMS classes will not be validity- or authority-checked.

Default: on an SMS system, for data sets which are SMS-managed and must be allocated, the SMS storage group ACS routine will be invoked to select a storage group and SMS will select a SMS-managed volume and allocate and catalog the data sets.

BYPASSACS and BYPASSSMS are primarily for use by storage administration personnel, since they bypass normal SMS allocation controls and rules. In order to use BYPASSACS or BYPASSSMS, the user of FDRCOPY must be authorized to the RACF profile

STGADMIN.ADR.COPY.BYPASSACS
in class FACILITY, or the equivalent in other security systems.

CATIFALLOC Specifies that **non-VSAM** output data sets will be cataloged even if they were preallocated (not allocated by the copy), subject to the NOCAT and RECAT operands (see [Section 21.01](#) "Cataloging non-VSAM data sets").

Default: output data sets are cataloged only when the FDRCOPY allocates them.

DATA= **USED** – only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be copied. On most volumes, this will make the copy run faster.

ALL – all allocated tracks of all data sets will be copied. You may need to specify DATA=ALL if the data sets to be copied include JES2 spool data sets or CICS log/journal data sets, since they usually do not have valid last block pointers.

Default: USED.

CONTINUED . . .

21.04 CONTINUED . . .

DSNENQ=

Specifies whether all of the input and output data sets selected for copy/move will be ENQed. See "Data Set Enqueue Option" in [Section 21.02](#) for more details. If the ENQ fails on an input data set, meaning that some other task has the data set enqueued, a FDR158 warning message is issued for the data set and it is bypassed unless the ENQERR=PROCESS operand is specified. An ENQ failure on an output data set will prevent the copy/move, unless FDRCOPY allocated the data set. A successful ENQ will prevent any other task from using the data set until the copy from that input volume is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be copied. The options for DSNENQ= are:

USE – The data sets will be enqueued for the duration of the copy from this disk volume. This is the most frequently used option and is the default for FDRCOPY.

TEST – The data sets will only be tested to see if they are enqueued to another task at the time that the copy from this volume starts.

HAVE – The data sets will be enqueued for the duration of the copy from this disk volume. If not available, a message (FDRW27) is issued to the MVS operator. See "Data Set Enqueue Option" in [section 21.02](#) for the valid responses.

NONE – No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: USE.

Recommendation: use DSNENQ=USE (the default) or HAVE if you want to be sure that no other task uses the data set until the copy/move is complete. However, use DSNENQ=NONE when another data set by the same name on another volume may be in use (e.g., copying data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

Note: a MOVE will delete the input data sets using IBM SCRATCH or DELETE services. These services may do their own SYSDSN ENQ on the data sets and may fail the deletion if the data set is in use.

ENQERR=

NO – If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see "Step Termination" in [Section 21.02](#)). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally.

Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.

ENQERR=

Specifies processing if the DSNENQ= option finds that an input data set is in use (enqueued):

BYPASS – do not copy an active data set.

PROCESS – copy a data set even if it is active (a warning message will still be produced).

Default: PROCESS.

NOTE: both ENQERR=NO and ENQERR=BYPASS/PROCESS may be specified on the same statement.

CONTINUED . . .

21.04 CONTINUED . . .

HFS=

QUIESCE invokes special processing when HFS (Hierarchical File System) data sets are copied. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to Unix System Services, so a "quiesce" call is issued to prevent updates to the HFS data set during the copy. If the quiesce fails and ENQERR=BYPASS was specified, the HFS data set will not be copied. See "HFS" in [Section 80.11](#) for details and security requirements.

Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for **all** data sets being backed up, not just HFS data sets. HFS=QUIESCE does not apply when moving HFS data sets; they must be dismounted before the MOVE.

Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table ([see Section 90](#)). If you use the default or cannot run with superuser status, you should unmount the HFS file system before the backup to be sure of getting a usable backup.

ICFCAT=

Applies to ICF VSAM files only. Specifies the source of the catalog name to be used if an output ICF VSAM cluster must be allocated.

ORIGINAL – use the catalog in which the input cluster is cataloged. When copying a cluster to a new name, ICFCAT=ORIGINAL is treated like ICFCAT=ALIAS, described below. If you need to catalog the output cluster into the same catalog as the input cluster but that catalog is not the one aliased for the new name, you must specify ICFCAT=STEPCAT and supply a STEPCAT DD statement pointing to that catalog.

STEPCAT – use the STEPCAT as the target catalog. If a STEPCAT DD statement is not supplied, it will use the master catalog or the catalog that is aliased for this data set in the master catalog.

ALIAS – determine the catalog from the alias name in the master catalog. If no alias is found and the cluster is being copied/moved to the same name, use the input cluster's original catalog. If no alias is found, and the cluster is being copied to a new name, FDRCOPY will use the STEPCAT (if present in the JCL) or the master catalog. Multi-level alias (MLA) is supported.

Default: ORIGINAL, except that if the cluster is being copied or moved to a newname (NEWGROUP or NEWINDEX specified) the default is ALIAS. If the output cluster is SMS-managed, ALIAS is forced.

NOTE: if you receive message FDR157 with COMP=0004 and CODE=00120 for a cluster, this means that its original catalog name does not exist on this system. Use ICFCAT=ALIAS or STEPCAT to place the cluster into a catalog which exists.

ICFCORE=

Specifies that the size of the table used to store the ICF VSAM cluster and component names. The copy must save all of the component names and their associated clusters which exist on the current input disk volume, in order to match up VTOC DSCBs (with the component name) to cluster names (for selection). nnnnnn is specified in bytes and must be large enough to contain all the VSAM names (depending on the length and number of names).

NOTE: Specifying ICFCORE= will increase the region requirement by the value specified. The default value uses one of the dump buffers and imposes no additional memory requirement.

Default: 53248, which will hold about 650 components.

CONTINUED . . .

21.04 CONTINUED . . .

MAXCARDS=	<p>Accept additional SELECT/EXCLUDE statements (over 250).</p> <p>Default: 250 statements.</p>
MAXERR=	<p>The number of input or output disk errors that are permitted by FDRCOPY prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Errors on the input and output disks are counted separately, but MAXERR= specifies the maximum for each counter.</p> <p>Default is 20 errors.</p>
NOCAT RECAT	<p>NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.</p> <p>RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.</p> <p>Default: For a COPY operation, FDRCOPY will catalog output non-VSAM data sets unless they are currently cataloged (even if to the input volume). For a MOVE operation, FDRCOPY will catalog output non-VSAM data sets unless they are currently cataloged to a volume other than the input. See Section 21.01 "Cataloging non-VSAM data sets" for more information.</p> <p>NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being copied and it is currently cataloged to another volume you can either specify RECAT or delete the output data set before the copy.</p> <p>NOCAT and RECAT are mutually exclusive. FDRCOPY will normally only attempt to catalog output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is specified.</p>
PRESTAGE	<p>Output data sets which already exist on the target output volume will not be copied. This may be used to avoid copying data sets which have already been copied. If the output data sets do not exist on the target volume, they will be allocated and copied.</p> <p>Default: pre-allocated data sets will have their contents overlaid.</p>
RLSE %FREE=	<p>RLSE – all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn – a percentage (nn%) of the PS and PO data sets to be left free after the copy. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the copy; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement).</p>

21.04 CONTINUED . . .

- SELTERR=** Specifies what will happen at step termination if FDRCOPY finds that a SELECT or EXCLUDE statement was never referenced (no data set on any input disk was selected by the statement):
- NO** – a condition code or ABEND is **not** to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.
- YES** – a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.
- Default: YES unless overridden in the FDR Global Option Table ([See Section 90](#)).
- SMSGDG=** Specifies the status of SMS-managed GDG (Generation Data Group) data sets, if allocated by FDRCOPY.
- DEFERRED, ACTIVE, or ROLLEDOFF** will set the GDG to that status.
- INPUT** will set the GDG to the SMS status of the input GDG generation. If the input GDG is non-SMS, it will be set ACTIVE if that generation is **currently** cataloged, otherwise DEFERRED.
- Default is DEFERRED, except for MOVE to the same name (no NEWNAME, NEWGROUP, or NEWINDEX parameter) when the default is INPUT.
- SNAP=** Used only when you are licensed for FDR InstantBackup and only if you have previously used program FDRSNAP to create an instant point-in-time image of a volume in an IBM RVA or StorageTek Iceberg/SVA disk subsystem with the Snapshot feature. [See Section 26](#) for details of the use of SNAP= with FDR.
- USE** - tells FDR InstantBackup to read the offline snapped copy of the online volume specified in JCL. FDR remembers the device address of the offline device most recently used as a target for a snap of each online volume unless an intervening IPL has occurred.
- (USE,REL)** - same as SNAP=USE, but at the end of copying from each snapped volume, FDR will issue a request to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the disk subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.
- Default: backup the online disk for each volume

21.05 FDRCOPY SELECT/EXCLUDE STATEMENT FOR COPY/MOVE

```

SELECT  DSN=filter
        S  CATDSN=filter
          DD=ddname
EXCLUDE ALLDSN
        X  FROM(CYL=cccc,TRK=tttt),TO(CYL=cccc,TRK=tttt)

          ,BLKF=nn                                ,NOCAT
                                                  ,RECAT

          ,CATALOG=catname
          ,MCATALOG=catname                      ,NVOL=(vvvvvv,vvvvvv,...)

          ,CATLIMITGDG=n                          ,PRESTAGE

          ,DATA=ALL                                ,PRTALIAS

          ,DATACLAS=dataclass                      ,RLSE
          ,NULLDATACLAS                          ,%FREE=nn

          ,DSNENQ=NONE                            ,STORCLAS=storageclass
                                                  ,NULLSTORCLAS

          ,DSORG=(xx,xx,...)

          ,MGMTCLAS=managementclass              ,TAPEDD=x
          ,NULLMGMTCLAS                          ,VOL=vvvvvv

                                                  ,TRK=nnnnnn
                                                  ,CYL=nnnnnn

          ,NEWNAME=newname
          ,NEWGROUP=newgroup
          ,NEWINDEX=newindex
          ,NEWDD=newddname

```

**SELECT/
EXCLUDE
STATEMENT
for COPY or
MOVE**

The SELECT statement selects the data sets which will be copied or moved. FDRCOPY will scan the VTOC of each input volume for data sets which match the parameters on a SELECT statements. Data sets may be selected by fully-qualified data set name or by using generic data set name selection (DSN=), or by copying a data set name from a DD statement (DD=). All data sets on an input volume may be selected (ALLDSN). If the SELECT statement specifies the VOL= operand, the named volumes are automatically processed as input volumes.

Data sets may also be selected from the system catalogs by fully-qualified name or using a generic filter (CATDSN=). The volumes found in the catalog entries will be automatically processed as input volumes, and the data set name found in the catalogs will be selected from their VTOCs. **Note:** the volumes processed due to CATDSN= will also be compared to other SELECT statements in the FDRCOPY step; if you have other SELECTs without either VOL= or TAPEDD= this may cause data sets to be selected unexpectedly.

It is also possible to copy specific tracks or ranges of tracks or cylinders to the identical locations on the output volume (FROM-TO). This operation will be a copy regardless of the COPY or MOVE statement, since these tracks will not be associated with any particular data set, and no data on the input disk volume will be disturbed. However, no check is made to see if the equivalent tracks being overlaid on the output disk volume belong to any particular data set, so this function should be used with care. Absolute track operations can be disabled by an option in the FDR Global Option Table.

The EXCLUDE statement prevents certain data sets from being copied or moved. The data sets to be excluded may be specified by fully-qualified name or by using generic data set name selection (DSN=) or copied from a DD statement (DD=). The EXCLUDE statement may be used to exclude particular data sets that would be selected by a more-encompassing SELECT statement. The EXCLUDE statement with the FROM/TO operands may be used to exclude particular tracks from the copy or move of a data set. **Since SELECT/EXCLUDE statements are scanned in the order they are input, EXCLUDE statements should usually precede SELECT statements.**

SELECT and EXCLUDE statements will apply to all input disk volumes unless TAPEDD= or VOL= parameters are specified.

21.05 CONTINUED . . .

**COPYING ICF
VSAM FILES**

FDRCOPY supports copying ICF VSAM clusters by cluster name using the DSN=, CATDSN= or ALLDSN operands. All components of a cluster on a given input volume will be copied. This includes alternate indexes and key range components. The component names will be reported by FDRCOPY followed by the base cluster name. All of the associated VVR information from the VVDS will be copied to the VVDS of the output disk. See "ICF VSAM Support" and the following topics in [Section 21.01](#) for other considerations for copying or moving VSAM.

OPERANDS DSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the VTOCs of selected input volumes.

EXAMPLES : DSN=USER1.JCL.CNTL
DSN=**LIST
DSN=PROD++.**.LIB*

DSN= does not have any special support for selecting GDGs, you can use CATDSN= to select relative generation number for a GDG.

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= parameter (if there are no DISKx DD statements pointing to those volumes, FDRCOPY will dynamically allocate and process them as input volumes). Specification of a relative generation number for GDG data sets is supported, e.g., CATDSN=A.B(-1)

If a filter is specified, then catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT CATDSN=dsname was present for each of them. It may be necessary to specify MAXCARDS=nnnnn if a large number of data sets are selected by the filter.

Additional considerations for CATDSN=filter are explained in [Section 80.14](#).

CATDSN= is supported only on SELECT statements. However, a preceding EXCLUDE statement with DSN= and/or VOL= can exclude data sets from selection by CATDSN=.

If the VOL= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

EXAMPLES : CATDSN=USER1.JCL.CNTL
CATDSN=**MASTER(0)
CATDSN=PROD++.**.LIB*

Normally CATDSN= will not display the data sets it selects from the catalogs, you will see the names only when FDRCOPY actually finds and selects the data sets in the VTOCs of the volumes they are cataloged to. To display all of the data sets selected specify PCATDSN=filter.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

CONTINUED . . .

21.05 CONTINUED . . .

DD= Specifies that a data set name is to be taken from a DD statement. The value of DD= is the DDNAME of a JCL statement. Using this option lets you specify a non-standard data set name or a generation data set (GDG) relative generation (CATDSN= will also let you select relative generations).

```
EXAMPLE: SELECT      DD=DD1
           //DD1      DD   DSN=A.B.C(0),DISP=SHR
```

Note that although DD= copies the data set name from the DD statement, it does not use the volume pointed to by that DD. The data set will be searched for only on the volumes being processed; the volume containing the data set must be among them. You could have a DISKx pointing to a data set and use DD=DISKx, so that the volume containing the data set will automatically be included.

ALLDSN Specifies that FDRCOPY is to process all the data sets on the volumes specified. DSN=** is equivalent to ALLDSN.

FROM/TO Identifies a range of track addresses for an absolute track copy. Absolute track statements can be mixed with SELECT statements for data sets. The FROM and TO operands must appear on the same input record, and cannot be continued. FROM/TO cannot be used on an EXCLUDE statement. Absolute track operations can be disabled by an option in the FDR Global Option Table ([Section 90](#)).

The cylinder (CYL=) and track (TRK=) addresses are in decimal, relative to zero. Leading zeros can be omitted. For example, valid specification on a 3390-3, which has 3339 cylinders and 15 tracks/cylinder, are CYL=0 TRK=0 to CYL=3338 TRK=14. The FROM address must be lower than or equal to the TO address.

If the FROM track (TRK=) is omitted, zero is assumed. If the TO track (TRK=) is omitted, the last track of the cylinder is assumed. So, if you specify only cylinder numbers (e.g., FROM(CYL=5),TO(CYL=7)) all tracks in those cylinders are processed.

NOTE: DSN=, CATDSN=, DD=, ALLDSN and FROM/TO are mutually exclusive. One and only one of these operands must be specified on each SELECT or EXCLUDE card.

BLKF= Selected PS and PO data sets are to be reblocked during the copy; see BLKF= in [Section 21.04](#) for details.

Default: data sets are not reblocked unless BLKF= was specified on the COPY/MOVE statement. The copy will fail if the input data set has blocks larger than the track size of the output disk.

CATALOG=
MCATALOG= Specifies the name of a user catalog (CATALOG=) or alternate master catalog (MCATALOG=) to search when CATDSN= is specified. See [Section 80](#) for details.

Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter.

CONTINUED . . .

21.05 CONTINUED . . .

CATLIMITGDG= May be used with CATDSN=filter to limit the selection of GDGs from the catalogs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:

n will cause only the most recently created "n" generations to be selected.

-n will cause only generation (-n) to be selected.

Default is that all the generations of selected GDGs will be selected (unless a relative generation number is specified at the end of the filter, e.g., CATDSN=filter(—2)).

DATA= **ALL** – specifies that FDRCOPY will copy/move the entire allocated space of the selected data sets. Normally it will only process up to the last block pointer (end-of-file) on input PS or PO data sets. Should be used if the last block pointer of certain data sets is invalid. DATA=ALL should not be used with RLSE or %FREE=.

DATACLAS=
NULLDATACLAS On a system with SMS active, specifies the SMS data class to be associated with the output data, overriding the original data class of the data set (if any). The Data Class ACS routine will not be invoked.

NULLDATACLAS changes the data class to null (not specified).

Default: the original data class of the input data set (if any) will be associated with the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is set.

DSNENQ= **NONE** – the data set enqueue will not be done for the selected data sets. This can be used to override the DSNENQ= operand on the COPY/MOVE statement for certain data sets that you know will probably be ENQed by another task.

Default: enqueue option is determined by the DSNENQ option specified on the COPY or MOVE statement.

DSORG= Specifies that this SELECT/EXCLUDE statement is to apply only to data sets whose data set organization matches one of the DSORGs specified. If more than one DSORG is specified, they must be enclosed in parentheses.

Valid DSORGs are:

DA	--	BDAM	PS	--	SEQUENTIAL
IS	--	ISAM	PO	--	PARTITIONED
UN	--	UNDEFINED (NONE)	EF	--	ICF VSAM
			UM	--	UNMOVABLE

MGMTCLAS=
NULLMGMTCLAS On a system with SMS active, specifies the SMS management class to be presented to the SMS Management Class ACS routine for the output data set, overriding the original management class of the input data set (if any). The ACS routine may accept or override this class.

NULLMGMTCLAS changes the management class to null (not specified).

Default: the original management class of the input data set (if any) will be passed to the ACS routine for the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is passed.

CONTINUED . . .

21.05 CONTINUED . . .

**NEWNAME=
NEWN=**

Copy or moves the selected data set with a new name. NEWNAME should only be used when the SELECT statement selects only a single data set. If the newname ends in a GDG relative generation number, e.g., NEWNAME=gdgname(-1), a LOCATE will be done to get the proper absolute generation number. If NEWNAME= is used for a ICF VSAM cluster, IBM defaults will be used to name the components.

**NEWGROUP=
NEWG=**

Copy or move the selected data sets using a new group name. The characters specified will replace the beginning of the input data set name. Care should be taken when periods are used that index levels are not incorrectly changed. The new names will be validated for conformance to IBM naming standards.

EXAMPLE : SELECT DSN=ABC** , NEWG=XYZ

Any data sets copied will be renamed to start with characters XYZ.

21.05 CONTINUED . . .

**NEWINDEX=
NEWI=**

Copy or move the selected data sets using a new name formed by adding or replacing one or more index levels in the original name; replacement index levels do not have to be the same length as the original indexes they replace. In the simplest case, each index level specified in NEWI is used in place of the corresponding index in the original name. Any remaining index levels at the end of the name are copied unchanged. This can easily be used to change the first indexes of the name.

For example, if the input data set is A.B.C.D,

NEWI=D	results in	D . B . C . D	(first index replaced)
NEWI=DD . E	results in	DD . E . C . D	(first 2 indexes replaced)

If a period is specified without any preceding characters, one original index level is copied from the input data set name to the output. This allows you to easily modify indexes in the middle of the name.

For example, if the input data set is A.B.C.D,

NEWI=. . E	results in	A . B . E . D	(third index replaced)
NEWI=FF . . . G	results in	FF . B . C . G	(first and fourth indexes replaced)

If + is specified before a new index level, that new index is inserted into the output data set name at that point. If ++ precedes the new index, it will be added to the end of the name. If - is specified, the next input index level will be dropped from (not copied to) the output name.

For example, if the input data set is A.B.C.D,

NEWI=+F	results in	F . A . B . C . D	(new first index added)
NEWI=. . +F	results in	A . B . F . C . D	(new third index added)
NEWI=++F	results in	A . B . C . D . F	(new last index added)
NEWI=. . -	results in	A . B . D	(third index dropped)
NEWI=Q . - . +E	results in	Q . C . E . D	(combination)

Note that, except for the ++ option, every period in the NEWI= mask corresponds to a period (one index level) in the original (input) data set name. The resulting new name will be checked to insure it meets IBM standards.

If the NEWI= value ends in a GDG relative generation number, e.g., NEWI=..NEWMAST(-2), that relative number will be added to the end of the newname, and a LOCATE done to get the proper absolute generation number.

NEWI= is a convenient way to rename every input data set, while using some index levels from the original name and replacing other indexes or adding new indexes.

NEWDD=

Specifies the name of a DD statement from which the new name of the output data set is obtained. NEWDD= should only be used when the SELECT statement selects only a single data set.

NOTE: NEWN=, NEWG=, NEWI=, and NEWDD= are mutually exclusive. If none of them are specified, the data set will be copied or moved under its original name. One of these parameters is required when copying ICF VSAM clusters since clusters cannot be copied under their original name; however a new name is not required when moving ICF VSAM clusters (see ICF VSAM in [Section 21.01](#)).

21.05 CONTINUED . . .

NOCAT
RECAT

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged, unless overridden by NOCAT/RECAT on the COPY/MOVE statement.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being copied and it is currently cataloged to volume other than the input volume, you can either specify RECAT or delete the data set before the copy.

NOCAT and RECAT are mutually exclusive. The copy will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified on the COPY/MOVE statement.

CONTINUED . . .

21.05 CONTINUED . . .

NVOL=

Specifies the volume serial(s) of output disk volumes to which data sets selected by this statement are to be copied/moved. You may specify:

- 1) A single specific volume serial,
NVOL=ABC123
- 2) A list of specific volume serials, enclosed in parentheses,
e.g., NVOL=(TSO001,TSO002,TSO003)
- 3) A group of volumes by placing an asterisk at the end of the volume serial prefix,
e.g., NVOL=TSO*
- 4) A combination of specific and group,
e.g., NVOL=(TSO*,PROD*,ABC001)
- 5) All online disk volumes may be selected by NVOL=*

A list of online target volumes matching your specification is generated by scanning all disk UCBs in the system UCB chains; there is no guarantee of the order in which UCBs are found, so you cannot predict the order of the volume serials in the list. If you specify volume serials or groups which are not online, they are ignored and no error message will result.

However, if the first or only specification is a specific volume serial, it will be chosen as the first target volume, with other volumes placed after it in UCB chain order.

Also, if you are restoring a multi-volume data set (non-VSAM or SMS-managed VSAM), the volume sequence number of the piece of the data set being allocated will be used to select a specification from your list. For example, if NVOL=(A,B,C), the second piece of the data set will go to volume B. If that specification is a group, the first volume in the UCB chain matching that group will be tried. If the allocation is unsuccessful (such as insufficient free space), then other volumes in the NVOL list will be tried as described above.

The first target volume is checked to see if an output data set already exists there. If so, it copies over the existing allocation (unless PRESTAGE was specified). If not, it attempts to allocate the output data set on that volume. If the allocation fails, it will be retried on successive volumes in the list until it succeeds or until it fails on 64 volumes. If the list contains several disk device types, "like" volumes (same type as the data set being copied) will be tried first, then unlike devices.

For multi-volume data sets, a target volume is bypassed if a piece of the data set already exists there but is not the right piece, so that it will not attempt to copy the third volume of a data set on top of the first volume. When it finds a target volume in the list that does not contain a piece of the data set, it will be allocated.

Specifying multiple volume serials or a volume group allows you to copy data sets in one pass even when no one volume has available space to contain them all; they will be spread across many of the target volumes.

Default: the output volume will be selected by rules defined in [Section 21.01](#).

On a system with SMS active, NVOL= may be ignored if the data set does not exist on the volume specified and the data set is SMS-managed (see STORCLAS=).

Note: if an allocation is attempted on several volumes from your NVOL list but it fails on all of them, the message printed will usually show the allocation failure codes from the **first** volume only; failure codes from other volumes are not displayed and may be different.

CONTINUED . . .

21.05 CONTINUED . . .

PRESTAGE	<p>Specifies that selected data sets will not be copied if the output data set already exists on the first target output volume. This may be used to avoid copying data sets which have already been copied.</p> <p>Default: copy pre-allocated data sets, overlaying the existing contents of those data sets, unless PRESTAGE was specified on the COPY/MOVE statement.</p>
PRTALIAS	<p>When used on a SELECT statement with CATDSN= will display all of the alias names and user catalogs that were searched. This is effective only when the data set name mask you provided forces CATDSN= to start in the master catalog and search one or more alias catalogs. If the non-mask characters at the beginning of the mask are sufficient to cause CATDSN= to begin its search in a user catalog, PRTALIAS is ignored.</p>
RLSE %FREE=	<p>RLSE – specifies that all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn – specifies a percentage (nn%) of the PS and PO data sets to be left free after the copy. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the copy; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement or by RLSE/%FREE= on the COPY/MOVE statement).</p>
STORCLAS= NULLSTORCLAS	<p>On a system with SMS active, specifies the SMS storage class to be presented to the SMS Storage Class ACS routine for the data set being copied, overriding the original storage class of the data set (if any). The ACS routine may accept or override this class.</p> <p>NULLSTORCLAS changes the storage class to null (not specified).</p> <p>Default: the original storage class of the input data set (if any) will be passed to the ACS routine for the output data set. For a non-SMS input data set, a null class is passed.</p> <p>If the Storage Class ACS routine assigns a storage class to this data set or accepts the class passed, the data set will be allocated as SMS-managed., and the SMS Storage Group ACS routine may be invoked to determine the actual target volume. If the Storage Class ACS routine returns a null (blank) storage class name, the data set will be allocated as non-SMS and the FDRCOPY rules listed in Section 21.01 for volume selection must select a non-SMS target volume.</p>
TAPEDD=	<p>x – specifies the same character as specified in a DISKx DD statement. If this operand is present, then this SELECT/EXCLUDE will only apply to data sets on the input disk volume specified by the DISKx DD statement.</p> <p>TAPEDD= might be used when multiple DISKx DD statements point to the same input volume but the data sets are to be output to different output volumes as specified by TAPEx DD statements or NVOL= operands.</p>

21.05 CONTINUED . . .

TRK=
CYL=

If the data set selected by this SELECT statement must be allocated, CYL= or TRK= specifies the number of cylinders or tracks to be allocated to the data set. On PS or PO files when DATA=ALL was not specified, this value should be at least equal to the used portion of the data set. On all other types of files and when DATA=ALL is specified, this value should be equal to or greater than the original size of the file. For ICF VSAM clusters, modifies the size of the base data component only. If the space is too small for the data being copied, the copy will automatically extend the file for non-VSAM.

Default: use the original size of the data set.

VOL=

Specifies the input disk volume serial(s) to which this SELECT/EXCLUDE statement is to apply. It may specify a single serial (e.g., VOL=ABC123) or a group of volumes all starting with the same prefix (e.g., VOL=ABC*). If there are online disk volumes matching the VOL= value that are not pointed to by DISKx DD statements in the FDRCOPY step, FDRCOPY will automatically dynamically allocate them and process them as input volumes. If specified on a SELECT CATDSN=, only data sets cataloged to the specified volumes are chosen.

Default: If neither TAPEDD= or VOL= are specified, the SELECT/EXCLUDE statement will apply to all input volumes and SELECT CATDSN= will select data sets regardless of the cataloged volume serial.

NOTE: TAPEDD= and VOL= are mutually exclusive. One and only one of these operands may be specified.

21.10 FDRCOPY DATA SET COPY EXAMPLES

The COPY operation will create a duplicate of the original data sets without modifying the originals in any way. The output data sets may have the same name as the originals, or may be modified to have new names (using NEWNAME=, NEWGROUP=, or NEWINDEX=). If copying under the original names, the catalog may be updated to point to the new copy (RECAT parameter). By default, data sets will not be copied if they are currently in use (allocated to another task).

**COPY NON-
VSAM DATA
SETS TO
NEWNAME**

COPY non-VSAM data sets from one input volume (339012) to a new volume, giving them new names, and cataloging the new names. This might be used to create test copies of existing data sets. In this example, TEST.FILE1 will be copied as TEST.FILE2 and all data sets beginning with PAYROLL will be copied with a new second level index of TEST inserted; only PAYROLL files which are sequential or partitioned will be copied. The input volume (339012) is specified by a DISK1 DD statement, but could also have been specified by a VOL=339012 operand on the SELECT statements.

```
//COPYFILS EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD UNIT=DISK,VOL=SER=339012,DISP=OLD
//SYSIN DD *
COPY TYPE=DSF
SELECT DSN=TEST.FILE1,NEWNAME=TEST.FILE2,NVOL=339045
SELECT DSN=PAYROLL** ,NEWI=.+TEST,DSORG=(PS,PO),NVOL=339037
```

**COPY ICF
VSAM
CLUSTERS TO
NEWNAME
USING THE
CATALOG**

COPY ICF VSAM data sets to a new volume, giving them new names. In this example, all cataloged clusters whose cluster name begins with ABC will be copied and will have the first index level of the cluster and all components changed to XYZABC. Any clusters that are multi-volume will be properly copied. MAXCARDS=1000 was specified in case more than 250 clusters are selected.

```
//COPYVSAM EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
COPY TYPE=DSF,MAXCARDS=1000
SELECT CATDSN=ABC.** ,NEWI=XYZABC,
DSORG=EF,NVOL=(PROD01,PROD02)
```

NOTE: Since ICF VSAM clusters must be cataloged when created, the same cluster or component name cannot be cataloged on two different volumes at the same time, so ICF VSAM files must be copied using a newname (the NEWINDEX= or NEWGROUP= parameter). However, ICF VSAM clusters may be moved under their original name (FDRCOPY uses a temporary cluster name for the output, then scratches the original and renames the output when the move is successfully completed). See ICF VSAM in [Section 21.01](#) for details.

21.11 FDRCOPY DATA SET MOVE EXAMPLES

The MOVE operation will move data sets from one disk volume to another. The output data sets may have the same name as the originals, or may be modified to have new names (using NEWNAME=, NEWGROUP=, or NEWINDEX=). In either case, if the move is successful, the original data set will be scratched and the catalog will be updated to point to the new data set. By default, data sets will not be moved if they are currently in use (allocated to another task).

MOVE NON-VSAM VOLUMES MOVE non-VSAM data sets to a new volume under their original names, and scratch the original data sets. They will be recataloged to the output volume unless they are currently cataloged to a volume other than the input volume. All ISPF data sets ("ISPF" anywhere in their names) on the input volume will be moved except PROFILE data sets (last index level of "PROFILE"). The recatalog and scratch will be done only if the copy was successful. The EXCLUDE statement will automatically exclude all VSAM files. **The input and output volumes are specified by DISK1 and TAPE1 DD statements, but could also have been specified by VOL= and NVOL= operands on the SELECT statement.**

```
//MOVEFILS EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD UNIT=DISK,VOL=SER=3380K1,DISP=OLD
//TAPE1 DD UNIT=DISK,VOL=SER=3380K2,DISP=OLD
//SYSIN DD *
MOVE TYPE=DSF
EXCLUDE ALLDSN,DSORG=EF
EXCLUDE DSN=**.*.PROFILE
SELECT DSN=**.*.ISPF**
```

MOVE DATA SETS SELECTED FROM THE CATALOG MOVE cataloged data sets to a new volume. The CATDSN= operand will cause the volume(s) the data sets are cataloged on to be processed as input volume(s). One specific data set, plus all DB2 data sets, will be moved (the mask on the second SELECT will select only DB2 clusters). All selected data sets will be moved to volume 3390T1.

```
//MOVEFILE EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
MOVE TYPE=DSF
SELECT CATDSN=MY.DATASET,NVOL=3390T1
SELECT CATDSN=*.DSNDBC.*.*.10001.A*,NVOL=3390T1
```

COMBINE SEVERAL VOLUMES ONTO ONE VOLUME Combine 3 3380 single density volumes (selected by VOL=3380J*) onto a 3380-K (triple density), by moving all of the data sets (including all non-VSAM and all single-volume ICF VSAM clusters). If the data sets were cataloged to the input volume the catalog will be updated to point to the data set's new volume. If the input volumes contain indexed VTOCs, VVDSs, or ABR models they will automatically be excluded (message FDR159 REASON=3 may be produced and can be ignored; no ABEND or error code will result). The 3 input volumes will be processed one at a time to avoid head movement on the output volume. Since the input volumes may contain a large number of data sets or ICF VSAM clusters which may increase FDRCOPY's storage requirements, the region requested is increased from the usual 2MB to 4MB.

```
//COMBINEK EXEC PGM=FDRCOPY,REGION=4M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
MOVE TYPE=DSF
SELECT ALLDSN,VOL=3380J*,NVOL=3380K1
```

21.11 CONTINUED . . .

**SPLIT ONE
VOLUME INTO
SEVERAL
VOLUMES**

MOVE all data sets from a 3390-3 (triple density) to several 3390-2 (double density) volumes (including all non-VSAM and all single-volume ICF VSAM clusters). If the data sets were cataloged to the input volume the catalog will be updated to point to the data set's new volumes. If the input volume contains an indexed VTOC, VVDS, or ABR model they will automatically be excluded (message FDR159 REASON=3 may be produced and can be ignored; no ABEND or error code will result). "NVOL=3390D*" will select all online disks whose volser starts with "3390D" for output; the disk with the lowest device address will be used until data sets can no longer be allocated on it, then the next such volume will be used.

```
//SPLITVOL EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
        MOVE TYPE=DSF
        SELECT ALLDSN,VOL=3390T1,NVOL=3390D*
```

21.12 FDRCOPY SMS EXAMPLES

The following are examples of using FDRCOPY to copy and move data sets on a system which has SMS (System Managed Storage) active. They illustrate techniques for converting data sets to and from SMS management. Output data sets will be SMS-managed if the SMS storage class ACS routine assigns a storage class to the data set or if it accepts the storage class passed to it by FDRCOPY. ACS routines are coded by each installation, so the decision on whether a data set is to be SMS-managed is a local one, and is usually out of the control of FDRCOPY (and the end-user).

When using FDRCOPY to create output data sets which are SMS-managed, you need not specify an output volume. If the output data set is SMS-managed, SMS will assign an output volume. However, if any output data set is not assigned a storage class, it will not be copied/moved unless a non-SMS target volume is specified by a TAPEX DD or by the NVOL= operand.

**CONVERT
DATA SETS TO
SMS**

MOVE a set of data sets from 3 non-SMS volumes, converting them to SMS management. The SMS ACS routines must be coded to assign these data sets a storage class; SMS will select an output volume for each and allocate the data sets. The original data sets will be scratched, and the output data sets will be cataloged on the SMS volumes.

```
//CONVSMS      EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
MOVE           TYPE=DSF
SELECT         DSN=PROD**,VOL=PROD01
SELECT         DSN=PROD**,VOL=PROD02
SELECT         DSN=PROD**,VOL=PROD03
```

**COPY SMS-
MANAGED
DATA SETS**

COPY a set of SMS-managed data sets, creating SMS-managed copies. The original storage and management classes of the input data sets will be passed to the SMS ACS routines, which may accept or override them. Data sets with a first-level index beginning with "APPL" will be copied. The NEWI= parameter is used to insert "TEST" as a new second-level index in all the output data set names.

```
//COPYSMS      EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
COPY           TYPE=DSF
SELECT         CATDSN=APPL*.**,NEWI=.+TEST
```

**OVERRIDE SMS
CLASSES**

MOVE a set of data sets (which may be SMS-managed or non-SMS), making them SMS-managed and requesting that certain SMS classes be assigned to the output data sets. The storage and management classes specified will be passed to the SMS ACS routines, which may accept or override them. The input data sets will be scratched and the output data sets will be cataloged on the output SMS volumes.

```
//MOVESMS      EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=DISK,VOL=SER=ABC123,DISP=OLD
//SYSIN        DD     *
MOVE           TYPE=DSF
SELECT         DSN=APPL2**,STORCLAS=APPL,MGMTCLAS=WKBACKUP
```

21.12 CONTINUED . . .

**CONVERT SMS
DATA SETS TO
NON-SMS**

MOVE all data sets off of an SMS-managed volume, requesting that they be converted to non-SMS. The SMS ACS routines will be invoked, but the storage class routine will be passed a null storage class; if it accepts the null class, FDRCOPY will allocate the output data sets as non-SMS on the volumes specified by NVOL=.

```
//MOVENSMS      EXEC   PGM=FDRCOPY , REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
MOVE           TYPE=DSF
SELECT        ALLDSN, NULLSTORCLAS, VOL=SMS123, NVOL=(XYZ234, XYZ345)
```

**BYPASS ACS
ROUTINES**

An authorized user (such as a storage administrator) may need to bypass the SMS ACS routines, to force a data set to be SMS-managed or non-SMS-managed, and to specify the SMS classes to be used. In this example, TEST.DATASET1 will be assigned the specified SMS classes; SMS will be invoked to select a storage group and assign a volume. TEST.DATASET2 will be assigned a null storage group, so it will be moved to the non-SMS volume specified by TAPE1. See the description of BYPASSACS in [Section 21.04](#) for authorization requirements.

```
//BYPASS        EXEC   PGM=FDRCOPY , REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//DISK1         DD     UNIT=DISK, VOL=SER=SMS123, DISP=OLD
//TAPE1         DD     UNIT=DISK, VOL=SER=XYZ234, DISP=OLD
//SYSIN         DD     *
MOVE           TYPE=DSF, BYPASSACS
SELECT        DSN=TEST.DATASET1, MGMTCLAS=TESTDS, STORCLAS=TEST
SELECT        DSN=TEST.DATASET2, NULLSTORCLAS
```

**BYPASS SMS
ALLOCATION**

An authorized user (such as a storage administrator) may need to force the allocation of SMS-managed data sets onto specific volumes; normal SMS facilities will not allow you to do so, but FDRCOPY does. Each data set selected will be passed to the SMS storage and management class ACS routines; each data set to which SMS assigns a storage class will be allocated and cataloged on the volume indicated by NVOL=; this volume MUST be a SMS-managed volume. Any data set which does not get a storage class assigned will not be moved. See the description of BYPASSSMS in [Section 21.04](#) for authorization requirements.

```
//BYPASS        EXEC   PGM=FDRCOPY , REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
MOVE           TYPE=DSF, BYPASSSMS
SELECT        CATDSN=DEVEL**, NVOL=SMS123
```

NOTE: BYPASSACS and BYPASSSMS may be used together. If so, any data set which was originally SMS-managed or which has a SMS storage class specified by STORCLAS= will be directly managed and cataloged as SMS-managed on the SMS-volume designated by TAPEX.

21.13 FDRCOPY UNLIKE DEVICE EXAMPLES

FDRCOPY can copy/move most types of data sets to a "unlike" disk device using a "logical" operation; for example, 3380 data sets can be moved to 3390 disks. Logical copy/move is automatically invoked when the output device type of a given data set differs from the type of the input disk. Special Considerations can be found in [Section 80](#), and in member UNLIKE in the FDR ICL (Installation Control Library). For ICF VSAM in particular, refer to member VSAMUNLK in the ICL for considerations on moving clusters to smaller disks (3390 to 3380). Most of the examples in the preceding sections will work even if the output disk type is different from the input type. Following are some specific unlike move examples.

**MOVE DATA
SETS TO 3390**

MOVE 3380 data sets to a 3390 disk. The data sets will be cataloged to the output volume. All data sets will retain their original block sizes; they will be allocated with a size about equal in bytes to the original 3380 data set. The selected data sets may include ICF VSAM clusters and non-VSAM data sets. Multi-volume VSAM clusters cannot be moved by this procedure, [See Section 80](#) for Multi-Volume VSAM Considerations.

```
//UNLIKE      EXEC  PGM=FDRCOPY,REGION=2M
//SYSPRINT    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//SYSIN       DD    *
      MOVE     TYPE=DSF,RECAT
      SELECT   CATDSN=TEST.*,NVOL=D3390X
/*
```

**MOVE DATA
SETS TO 3380**

MOVE PS and PO data sets from a 3390 disk to several 3380 disks (volsters starting with "D3380"). The data sets will be cataloged to their output volume. All PS (sequential) data sets will be reblocked to half-track blocking for maximum track utilization.

```
//UNLIKE      EXEC  PGM=FDRCOPY,REGION=2M
//SYSPRINT    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//SYSIN       DD    *
      MOVE     TYPE=DSF,RECAT,BLKFB=2
      SELECT   DSN=**TEST**,VOL=D3390A,NVOL=D3380*,DSORG=(PS,PO)
/*
```

**MOVE ENTIRE
3380 TO 3390**

MOVE all data sets from a 3380 disk to several 3390 disks. In most cases, all the data sets on a 3380 fit on a 3390-2, but multiple 3390 output volumes are provided in case there is not room. All data sets will be cataloged to their output volume. All PS (sequential) data sets will be reblocked to half-track blocking for maximum track utilization; PO data sets will have their block size increased to half-track for new members. The selected data sets may include ICF VSAM clusters and non-VSAM data sets. Multi-volume VSAM clusters cannot be moved by this procedure, [See Section 80](#) for Multi-Volume VSAM Considerations.

```
//UNLIKE      EXEC  PGM=FDRCOPY,REGION=2M
//SYSPRINT    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//SYSIN       DD    *
      MOVE     TYPE=DSF,RECAT,BLKFB=2
      SELECT   ALLDSN,VOL=D3380K,NVOL=(D3390A,D3390B,D3390C)
/*
```

21.14 FDR INSTANTBACKUP EXAMPLES

This section contains an example of the use of FDRCOPY with FDR InstantBackup. FDR InstantBackup is described in more detail in [Sections 25 through 29](#), depending on the type of disk subsystem you are using; there are more examples for each subsystem in those sections.

FDR InstantBackup is an additional cost enhancement to FDR.

NOTE: FDR InstantBackup enhances FDRCOPY to use the hardware features of certain disk subsystems to quickly move data tracks. Currently this supports the IBM RVA or StorageTek Iceberg/SVA with the Snapshot feature and the EMC Symmetrix with the Timefinder feature. No special parameters are required; when FDRCOPY detects that the input and output disk volumes for a given data set being copied are in the same disk subsystem with the proper features, those features are invoked and the copy/move will run in significantly less elapsed time. All of the examples in the preceding sections will automatically use FDR InstantBackup if the proper hardware is involved. You can tell if this was done by looking for the word INSTANT in the FDR311 message for each data set.

The example below demonstrates another FDR InstantBackup enhancement for FDRCOPY. They show how you can "freeze" a point-in-time copy of an online disk volume to another offline volume, and then copy data sets from that frozen offline copy.

**INSTANT
COPY FROM
EMC
SYMMETRIX**

Data sets are to be copied from an EMC Symmetrix with the Timefinder feature. The BCV at address 01FA has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous one-time ESTABLISH has been issued to establish that pairing. The step SPLIT will split the BCV from its online volume and wait for the split to complete, creating a point-in-time image of the volume. Step COPY will backup the BCV copy of data sets from volume PROD01 to newnames on a new volume. After the copy is complete, step REEST will re-synchronize the BCV with the online volume. [See Section 25](#) for details of FDR InstantBackup for EMC Symmetrix.

```
//SPLIT EXEC PGM=EMCTF
//SYSOUT DD SYSOUT=*
//SYSIN DD *
    SPLIT 1,01FA,WAIT
//COPY EXEC PGM=FDRCOPY,REGION=0M,COND=(0,NE,SPLIT)
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR.USE.UNIT01FA,UNIT=SYSALLDA,
// VOL=SER=PROD01,DISP=OLD
//SYSIN DD *
    COPY TYPE=DSF,BCV=(USE,RET)
    SELECT DSN=LEDGER**,NEWI=LEDTST,NVOL=TEST*
```

Details of FDR InstantBackup for supported disk systems are found in [Sections 25-29](#).

21.20 FDRCOPY ISPF INTERFACE

For ABR customers, FDRCOPY is supported by the ABR SRS dialogs on ISPF. As described in [Section 54.50](#), SRS allows end users to generate a table of selected data sets and their characteristics, and to enter commands to be executed against those data sets.

COPY PANEL If the "copy" command is entered for a data set, you will receive a panel similar to this:

```
----- FDRSRS - Copy Data Set -----
COMMAND ==>                                SCROLL ==> HALF
      Edit generated JCL      Submit generated JCL      FG - execute in the foreground
Operands for COPY TYPE=DSF statement (section 21.04):
==> COPY TYPE=DSF
FROM DSNNAME / Filter ==> 'BAB.TEST.LOAD'
      Volume Serial      ==> TS0001
TO   New DSNNAME         ==>
      or NEWINDEX        ==> .TEST2
New Volume Serial(s) ==> TS0*
Operands for SELECT DSN= statement (section 21.05):
==> NOTIFY=BAB2
```

The name and volser of the input data set will be filled in. Since this is a COPY, you will probably want to enter a new name for the output data set and output volumes, as shown above.

MOVE PANEL If the "move" command is entered for a data set, you will receive a panel similar to this:

```
----- FDRSRS - Move Data Set -----
COMMAND ==>                                SCROLL ==> HALF
      Edit generated JCL      Submit generated JCL      FG - execute in the foreground
Operands for MOVE TYPE=DSF statement (section 21.04):
==> MOVE TYPE=DSF
FROM DSNNAME / Filter ==> 'BAB.ISPF4.PROFILE'
      Volume Serial      ==> TS0001
TO   New DSNNAME         ==>
      or NEWINDEX        ==>
New Volume Serial(s) ==> SMS1*
Operands for SELECT DSN= statement (section 21.05):
==> NOTIFY=BAB2
```

The name and volser of the input data set will be filled in. Since this is a MOVE, you will probably want to enter new volumes for the output data set, as shown above.

PANEL EXECUTION If needed, you can provide additional operands for the COPY/MOVE statement or the SELECT statement on the lines indicated.

Once you are done, you can submit a batch FDRCOPY jobstream to perform the copy/move (EDIT or SUBMIT commands) or to attach FDRCOPY under TSO and execute the copy/move in the foreground (FG). The foreground operation will tie up your TSO session until the operation is done.

25.01 FDR INSTANTBACKUP FOR EMC TIMEFINDER™

FDR InstantBackup is a additional cost FDR facility for non-disruptive backup. It works in concert with TimeFinder™, an EMC Symmetrix feature which allows for the "instant" creation of exact duplicates of existing DASD volumes.

FDR InstantBackup enhances:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY data set copy
- FASTCPK volume reorganization
- FDRABR volume backup (full-volume and incremental backup)

It provides non-disruptive backup of offline MVS DASD volumes, both SMS and non-SMS, using technology developed by Innovation Data Processing.

TIMEFINDER™

Traditional backups may disrupt normal operations. It obviously takes time to backup a data set or volume. If data sets are being updated by some application while they are being backed up, the backup may not be valid and the data sets may not be usable when they are restored. The usual response to this problem is to quiesce all updates to the data during the backup, disrupting normal operation until the backup is complete.

When used with EMC's TimeFinder feature, FDR InstantBackup allows you to instantly freeze a copy of an online disk volume and easily backup or copy data from that frozen copy as if it was coming from the original volume. Updates to the data need to be quiesced only for the few moments necessary to create the frozen copy.

**FDR
INSTANTBACKUP
ENHANCES FDR
COMPONENTS**

FDR InstantBackup enhances FDR, FDRDSF, and FDRCOPY to read from an offline copy of an online volume which was created as a point-in-time frozen image using Timefinder. This creates a backup or copy of the volumes or data sets that looks exactly like a normal backup or copy, except that the data is frozen in time.

FDR InstantBackup also enhances FDRCOPY to use an internal track copy feature available on many Symmetrix subsystems. When copying or moving data sets with FDRCOPY, if the input volume and the output volume for a given data set are both within the same Symmetrix, FDRCOPY will issue special commands to copy the required tracks internally, without sending the data over the channel and back again. This will run 20-50% faster than a normal copy. If the input and output volumes for a data set are not in the same Symmetrix or the copy feature is not available, normal reads and writes are done to copy the data. This feature is automatic; no special options are required to invoke it.

If you are licensed for COMPAKTOR, FDR InstantBackup also enhances FASTCPK to use that facility of EMC Symmetrix systems to move tracks internally without sending the track contents to the host and back again (READ and WRITE). If you specify EMCCOPY=YES on the COMPAKT statement, FASTCPK will use the Symmetrix facility to move tracks on the volume being COMPAKTed, resulting in a significant savings in elapsed time.

If you are licensed for FDRABR, FDR InstantBackup enhances ABR Volume Backups, both full-volume (TYPE=FDR) and incremental (TYPE=ABR/DSF/AUTO) to use Timefinder, by creating "instant" frozen copies of the volumes to be backed up, and then moving that captured data to tape.

CONTINUED . . .

25.02 TIMEFINDER OPERATION

**SYMMETRIX
DATA
PROTECTION
OPTIONS**

EMC Symmetrix systems allow you to assign logical disks within the Symmetrix as "mirrors" of online disks. These mirrors are automatically maintained as exact duplicates of their online volumes and all writes to the online volume are "mirrored" to the mirror disks. These mirrors are used for recovery from disk errors and hardware failures. A normal mirror disk does not have its own device address and is hidden from the operating system. Assignment of mirrors is done by EMC during the hardware configuration.

Normal mirrors are often called RAID-1 since they provide complete duplication of the online volume. Symmetrix systems also support other types of RAID protection of data.

**TIMEFINDER™
BUSINESS
CONTINU-
ANCE
VOLUMES**

TimeFinder™ is an optional Symmetrix facility. TimeFinder and FDR InstantBackup, working together, allow "instant" backups of volumes by "freezing" an image of a volume at a particular point in time, then backing up that frozen image. You can also copy data sets from the frozen image to other online disk volumes.

TimeFinder uses a new type of mirror, called a Business Continuance Volume or BCV. A given online volume may have normal mirrors as well as BCVs, but unlike normal mirrors, a BCV

- has its own device address
- is not permanently assigned to any online volume
- can be assigned (established) by host software as a BCV mirror of any online volume in the Symmetrix subsystem with the same size and format (3390 or 3380). While assigned, the BCV will always be an exact copy of its online volume
- can be detached (split) by host software from its currently assigned online volume
- after the split is a frozen point-in-time image of the online volume.

BCVs do not replace normal RAID-1 mirrors or RAID protection, since BCVs are not for data protection. A given online volume can be assigned a BCV regardless of the data protection scheme in use for that volume.

Your EMC representative can provide you with more information on TimeFinder and BCVs, including more detail on the internal operation of BCVs. EMC provides a utility program, EMCTF, for the management of BCVs. Examples of EMCTF control statements are shown in this section, but you should review the EMCTF documentation provided by EMC to verify control statement format and options.

**TIMEFINDER
WITHOUT FDR
INSTANTBACKUP**

Although TimeFinder can instantly create a point-in-time image of a production disk volume, traditional backup products cannot directly use the BCV image without additional considerations.

For example, a conventional backup product will require that the split BCV be brought online. Since MVS will not let you put two volumes with the same volume serial online, the BCV volume serial must be changed.

EMCTF provides a facility for re-labeling a BCV during the split process so that it can be put online, but re-labeling a volume invokes additional requirements for SMS managed volumes, VSAM clusters and cataloged data sets. The data sets on the re-labeled volume will appear to be uncataloged, which may cause problems during backup and restore unless they are renamed and re-cataloged.

A backup of a re-labeled volume will appear to be a backup from the new volume serial, not the original volume serial, so additional procedures are required to document this correspondence (e.g., backup of volume B is really a backup of volume A). Restore procedures must be modified to restore the backups to the correct serial (e.g., restore volume B back to volume A and re-label it as A).

EMC also provides a utility to rename and recatalog the data sets on the BCV after the split. Although this could be used to backup those data sets with a standard backup utility, the installation must now use special restore procedures to restore each data set from its temporary name on the backup to its true name on the online volume. For large numbers of data sets, this could be onerous.

25.03 THE FDR INSTANTBACKUP SOLUTION

FDR InstantBackup allows you to create and use a duplicate BCV volume image without changing its volume serial or bringing it online. FDR full volume backup, FDRDSF data set backup, FDRCOPY data set copy, and FDRABR volume backups use FDR InstantBackup technology to read the split BCV even while it remains offline. A non-disruptive backup or copy can start as soon as the BCV volume is split. Updates to the online volumes can continue while you create the point-in-time backup or copy of the data.

For FDRABR, FDR InstantBackup makes volume backups a 2-step process. The first step executes a SPLIT statement under FDRABR, to split the BCVs of the selected volumes. If the BCV is fully synchronized with its online volume, this takes only seconds per volume. The split BCV is actually a valid backup at that point-in-time; you can do a restore from it (using Timefinder facilities) if recovery is required. The second step executes a normal FDRABR full-volume or incremental backup of the same volumes, except that ABR will read the offline point-in-time image and move it to a backup on tape. There is no need to identify the BCV device to ABR; it knows the BCV address of any given online volume.

You invoke FDR InstantBackup for FDR, FDRDSF, and FDRCOPY with standard FDR JCL which points to the online volume but which contains a special data set name that identifies the offline BCV just split from the online volume. FDR InstantBackup verifies that the BCV device is offline and that its volume serial matches the serial of the online disk. It then will read the offline BCV in exactly the same way that it would read the online disk. FDR InstantBackup produces a backup that appears to be a conventional backup of the online volume serial. FDR InstantBackup blends its unique off-line, high-speed non-disruptive backup together with a traditional restore complete with FDR's powerful logical file capability.

There are no special considerations for restores from a backup created by FDR InstantBackup. You can restore entire volumes or individual data sets from those backups. The target volumes for restore will be online volumes, not BCVs. Since the backups created by FDR InstantBackup appear to be backups of the online volume, there are no special volume serial or data set name concerns.

FDRCOPY data set copy and move can be used with FDR InstantBackup in two ways:

- when FDRCOPY detects that the input and output volume associated with a given data set are in the same Symmetrix, it automatically invokes an internal EMC track copy facility to copy the tracks, resulting in a faster copy. Normal read/write is used if the volumes are not both in the same Symmetrix.
- Alternately, you can use Timefinder to freeze an image of the volumes containing the desired data sets, and then use FDRCOPY to create copies of those data sets on other online volumes. This is useful if the target volumes are not in the same Symmetrix as the source volumes, allowing you to capture an instant image of the source data sets and copy them at your leisure.

Note: Although Timefinder supports the Symmetrix SRDF feature, allowing the BCV to be in a different Symmetrix subsystem from the online volume, FDR InstantBackup only supports BCVs which are in the same subsystem. Also, both the online volume and its BCV copy must be accessible from the MVS system where FDR InstantBackup is run.

25.04 SETTING UP BCVS

**BCV
DEFINITION**

To define BCVs, you must have the TimeFinder support installed on your Symmetrix, and you must have sufficient disk capacity within the subsystem to define as BCVs. You must also install the EMC Timefinder utility, EMCTF.

To determine the number of BCVs to be defined, you must identify all of the online volumes for which you intend to use FDR InstantBackup and categorize them by type (3380 or 3390) and size (number of cylinders). A BCV can be assigned only to an online volume of identical type and size.

If you have sufficient capacity, you can define a BCV for every online volume involved. If not, you can use the techniques described later in this section to establish a "pool" of BCVs which can be assigned to online volumes as needed.

Your EMC representative will assign logical volumes within the Symmetrix as BCVs during the hardware configuration. Initially, the BCVs will not be associated with any online volume.

The device addresses assigned to the BCVs must be defined as normal 3380 or 3390 disks in the Symmetrix subsystem in your I/O configuration (IOCP or HCD). However, as shown later in this section, when a BCV is detached from its online volume for use with FDR InstantBackup it will have the same internal volume serial as the online volume. This is normally not a problem since it will be offline to MVS, but it can cause confusion if you IPL your MVS system while BCVs are split; IPL will detect that two volumes have the same serial and ask the operator which should be left offline (message IEA213A). Obviously, the real online volume should remain online so the operator must respond with the BCV address. **To avoid the possibility of the operator replying incorrectly, you may want to mark the BCVs as OFFLINE in your I/O configuration, so that MVS will not attempt to access them at IPL time.** FDR InstantBackup will be able access them even though they are offline. However, if you use the BCVs only with FDR InstantBackup for ABR, this is not necessary since ABR will modify the volume label to prevent the BCV from being varied online.

For example, in HCD, the MVS options for a BCV device will look similar to:

Command ==> View Device Parameter / Feature Definition				Row 1 of 6
				Scroll ==> PAGE
Configuration ID . . . : CURP1099		Common Configuration data		
Device number . . . : 01E0		Device type . . . : 3390		
Generic / VM device type : 3390				
Parameter/	Value	Req.	Description	
Feature				
OFFLINE	Yes		Device considered online or offline at IPL	
DYNAMIC	Yes		Device supports dynamic configuration	
LOCANY	Yes		UCB can reside in 31 bit storage	
ALTCTRL	No		Separate physical control unit path	
SHARED	Yes		Device shared with other systems	
SHAREDUP	No		Shared when system physically partitioned	

In IOCP/MVSCP, specify OFFLINE=YES on the IODEVICE macros for the BCV devices. The following is an example of a Symmetrix attached via 2 ESCON channels (CHPIDs 40 and 41), with 128 online disk addresses (300-37F) and 64 BCV addresses (380-3BF):

```

CNTLUNIT CUNUMBR=008,PATH=(40),LINK=(C0),
        UNIT=3990,UNITADD=((00,192))
CNTLUNIT CUNUMBR=009,PATH=(41),LINK=(C4),
        UNIT=3990,UNITADD=((00,192))
IODEVICE ADDRESS=(300,128),CUNUMBR=(008,009),  ** ONLINE
        UNIT=3390,UNITADD=00,FEATURE=SHARED
IODEVICE ADDRESS=(380,64),CUNUMBR=(008,009),  ** BCV
        UNIT=3390,UNITADD=80,FEATURE=SHARED,OFFLINE=YES

```

Note that some IOCP documentation may indicate that UNITADD=00 is required for ESCON; what it really means is that the value for the **first** IODEVICE definition for a given control unit must be 00. So when varying characteristics must be specified for some devices on a control unit (such as OFFLINE=YES), you can have multiple IODEVICE statements with varying values for UNITADD.

25.04 CONTINUED

**TIMEFINDER
WITH FDR
INSTANTBACKUP**

A BCV is first assigned as a mirror by an EMCTF ESTABLISH command. ESTABLISH will synchronize the BCV with the online volume by copying all data on the online volume to the BCV; this will take several minutes per Gigabyte (GB) to be copied until the volumes are synchronized. Once the copy is complete, all updates to the online volume will be mirrored on the BCV, **so the Symmetrix hardware will insure that the BCV is always an exact copy of its online volume.** While the BCV is assigned, the host cannot directly address it; it will appear to be "not ready".

Here is an example of an ESTABLISH command, where "01E4" is the online device and "01FA" is the BCV:

```
ESTABLISH 1,01FA,01E4
```

Since **all** data on the online volume is mirrored on the BCV, including the volume label and the VTOC, VTOCIX, and VVDS, the BCV also has the same volume serial as the online volume.

At the point in time where a backup or copy of the data on the online volume is desired, the BCV can be detached from the online volume by a SPLIT command. The BCV will be an exact copy of the online volume at the time that the SPLIT was issued. SPLIT takes very little time if the BCV is already synchronized. You may want to quiesce updates to the online volume for the few moments required to SPLIT.

For FDRABR, you split the BCV using the SPLIT statement of FDRABR. This allows ABR to identify the address of the BCV for later processing. See the examples later in this section for more details.

For FDR, FDRDSF, and FDRCOPY, you split the BCV using the EMCTF SPLIT command. Here is an example of a SPLIT command, where "01FA" is the BCV device (the Symmetrix knows to which online volume it is currently assigned):

```
SPLIT 1,01FA,WAIT
```

The WAIT operand waits for the SPLIT to complete so that you can be sure that the BCV is available before starting the backup or copy. EMCTF includes an option (VOLID) to relabel the BCV during the SPLIT and a command (RELABEL) to do so later, **but you must not use VOLID or RELABEL with FDR InstantBackup.**

Once split, the BCV can be accessed by its assigned device address and can be backed up or copied with FDR InstantBackup. FDR InstantBackup can backup or copy data from the BCV even though it is offline and has a serial that duplicates the online volume.

When the backup or copy is complete, you can re-synchronize the BCV with the online volume with a RE-ESTABLISH operation. During the time when the BCV was split, the Symmetrix keeps track of all data updated on the online volume and the BCV. RE-ESTABLISH copies only updated data from the online volume to the BCV, so it is much faster than a full ESTABLISH. Once the BCV is again synchronized with the online volume, it is ready to be used for another FDR InstantBackup operation.

For FDRABR, the ABR step which moves the backup data to tape has an option BCV=(USE,RET) which will automatically re-establish the BCV when the backup is complete.

For FDR, FDRDSF and FDRCOPY, you re-establish the BCV using the EMCTF RE-ESTABLISH command. Here is an example of a RE-ESTABLISH command, where "01FA" is the BCV device (the Symmetrix remembers the online volume from which it was last split):

```
RE-ESTABLISH 1,01FA
```

Note that the copying of data during the ESTABLISH and RE-ESTABLISH is done internally by the Symmetrix subsystem and is a low priority task (real I/Os get priority). However, the amount of time required is not really important unless you need to SPLIT the BCV and do a backup before the volumes are completely synchronized.

You never need to bring a BCV online or re-label it or re-initialize it with ICKDSF, as long as it is used only for FDR InstantBackup.

CONTINUED . . .

25.04 CONTINUED

**SETTING UP A BCV
ENVIRONMENT**

There are 2 ways to setup BCV use:

- 1) you can permanently assign BCVs to every online volume for which you plan to use BCV backup or copy. This will require more Symmetrix capacity and more device addresses, but is much simpler to administer. The BCV is actively mirroring it's online volume all the time, except during the backup or copy, so the ESTABLISH is done only once. A SPLIT is done before the backup or copy. Afterwards, the RE-ESTABLISH will quickly re-synchronize the volumes since only changed data is copied.
- 2) you can create a pool of BCVs and assign them to online volumes as necessary. This requires less Symmetrix capacity and fewer device addresses, but is more complex to administer. You will have to ESTABLISH the BCV every time it is to be used. It will take additional time to copy all of the data from the online volume, especially if multiple volumes are being ESTABLISHED. You will require operational procedures to insure that two FDR InstantBackup jobs don't try to use the same BCV volume for two different online volumes at the same time. A SPLIT is done before the backup or copy, but there is no need to RE-ESTABLISH the BCV afterwards.

You can use a combination of the two, some online volumes with permanent BCVs (perhaps for volumes with frequent or time-critical backups) and some with BCVs assigned from a pool (perhaps for less frequent or non-time-critical backups).

In either case, the BCV volumes must match the size (number of cylinders) and format (3380 or 3390) of the online volumes to which they will be assigned.

CONTINUED . . .

25.05 FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY

FDR InstantBackup uses standard FDR, FDRDSF and FDRCOPY JCL with one modification (see "FDR INSTANT BACKUP FOR ABR" in [Section 25.06](#) for FDRABR requirements). The FDR control statements used are standard except for a few considerations. For complete details on other JCL and control statement requirements of FDR, FDRDSF, and FDRCOPY, please see [Sections 10, 20, and 21](#) of the FDR manual, respectively.

FDR InstantBackup uses a special convention for the JCL statements that identify the disk to be backed up or copied, i.e. DISKx DDs. As shown in the examples below, the DD points to the online volume but includes a special data set name which tells FDR the device address of the BCV which is to be read in place of that online volume. FDR will verify that the volume serial of the online volume is the same as that of the offline BCV.

These special DISKx DDs can be used with:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY input disks for COPY operations only. Although FDRCOPY can copy data sets without the use of DISKx DD statements, DISKx DDs are required to cause FDRCOPY to read the offline BCV. You can still select data sets normally, including selecting them from the catalog, but FDR InstantBackup will be used for any input volume identified on a DISKx DD statement.

They can also be used for FDR full-volume copies (COPY TYPE=FDR) but in most cases it would be more useful to copy the online volume directly.

FDRCOPY MOVE operations from a point-in-time copy are not supported by FDR InstantBackup, since it deletes the input data set. You should always move the live, online version of a data set.

**NON-SMS
MANAGED
VOLUMES**

For non-SMS volumes, a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,UNIT=3390,
//          VOL=SER=vvvvvv,DISP=OLD
```

The UNIT and VOLSER (vvvvvv) must identify the online disk volume. Note that the FDR.USE.UNITuuuu data set does not really exist anywhere in your system, but the presence of the special data set name on the DD statement is very important. It is this name that directs FDR to access the offline BCV device located at UCB address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address. A leading zero must be added if it is a 3-digit address, e.g., DSN=FDR.USE.UNIT01F3. This special data set name is never opened so it will cause no problems with security.

CONTINUED . . .

25.05 CONTINUED

**SMS-MANAGED
VOLUMES**

For SMS-managed volumes a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,DISP=OLD
```

Because MVS will ignore a volume serial in the JCL if it points to a SMS-managed volume, the special FDR.USE.UNITuuuu name must be cataloged. **This data set does not really have to exist on the original volume, but it must be cataloged to that volume!** Just as for a non-SMS volume, its presence on the DD statement is very important. It is this name that directs FDR to access the BCV on the offline device located at device address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address; add a leading zero if necessary.

Here is one technique for creating the catalog entry required for FDR InstantBackup. This step will catalog the special FDR.USE.UNITuuuu name to an SMS managed volume. It could be run immediately before a backup.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE FDR.USE.UNITuuuu NOSCRATCH
DEFINE NONVSAM(NAME(FDR.USE.UNITuuuu) -
DEVICETYPE(33xx) VOLUME(vvvvvv))
```

The DEVICETYPE and VOLUME parameters identify the SMS-managed online volume. The special FDR.USE.UNITuuuu name will direct FDR to access the offline BCV device located at device address uuuu instead of the original online SMS managed online volume.

**FDR CONTROL
STATEMENTS**

The control statements used with FDR InstantBackup are the normal statements documented for FDR ([Section 10](#)), FDRDSF ([Section 20](#)), and FDRCOPY ([Section 21](#)) with these modifications:

- Since FDR InstantBackup is not reading the live data, it is never appropriate to use the ENQ= or DSNENQ= operands to enqueue on the VTOC of the BCV or the data sets on the BCV. To do so would unnecessarily prevent access to the live data on the online volume.
- For FDRCOPY, specify DSNENQ=NONE. Also, the input volumes must be specified by DISKx DD statements (to identify the BCV addresses) so you should not select data sets with CATDSN= or VOL=.
- Do not use the FDRCOPY MOVE statement when the input is a BCV. If you do so, the live data sets on the online volume will be deleted at the end of the move.
- If you wish to do a RE-ESTABLISH of the BCV with its online source volume as soon as each backup is done, add the parameter BCV=(USE,RET) to the DUMP statement. If you do so, a separate RE-ESTABLISH step is not required. However, all of the volumes processed in the step must be BCVs.

CONTINUED . . .

25.05 CONTINUED

**FDR
INSTANTBACKUP
OPERATION**

If you have a permanently assigned BCV volume for every online volume, use these steps:

- 1) Execute an EMCTF ESTABLISH one time for each online-BCV pair. There is no need to wait for the synchronization to complete.
- 2) When ready to backup or copy data, quiesce system and/or application update activity on the online volumes, if required, and execute an EMCTF SPLIT with the WAIT operand for each associated BCV. Once the SPLIT ends, you can re-enable updates on the online volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with the special DISKx DD statements described above to access the BCV.
- 4) When you no longer need the "frozen" data, execute an EMCTF RE-ESTABLISH for each BCV to resynchronize it with its online volume. Unless there has been a lot of update activity on the online volume, this will complete quickly.

If you are using a BCV pool (described earlier), use these steps:

- 1) Execute an EMCTF ESTABLISH for each online volume to be backed up, attaching it to one BCV from the pool. Be sure to do the ESTABLISH far enough in advance of the required backup time to allow the synchronization to complete. If you need to know when synchronization is complete (e.g., to know when you are ready to do the backups), use the WAIT operand.
- 2) When ready to backup or copy data, quiesce update activity on the online volumes, if required, and execute an EMCTF SPLIT with the WAIT operand for each associated BCV. Once the SPLIT ends, you can re-enable updates on the online volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with the special DISKx DD statements described above to access the BCV.
- 4) There is no need to RE-ESTABLISH the BCV unless you plan to do another BCV backup of the same online volume. Once the backup is complete the BCV is available for use with a different online volume. However, you will need operational procedures to insure that different FDR InstantBackup jobs do not attempt to use the same BCV at the same time.

**MULTIPLE
SYSTEMS**

FDR InstantBackup requires that both the original volume and its offline BCV copy be accessible from the MVS system where the backup is run. If you execute the SPLIT on one system, but execute the backup on a second MVS system which does NOT have access to the online volume, FDR InstantBackup will fail. There is a circumvention for FDR and FDRDSF, but you must contact Innovation for assistance in implementing it.

CONTINUED . . .

25.06 FDR INSTANTBACKUP FOR ABR

When a normal ABR full-volume or incremental backup completes successfully, it is recorded by doing 2 things:

- ABR catalogs the backup data set it just created in the ABR catalog
- ABR updates the VTOC on the volume to record in the ABR Model DSCB the generation and cycle of the backup just created, and updates the current backup information for every selected data set

So, as soon as the backup data set is complete and available, all ABR records point to it and restores can be immediately requested.

**OPERATION OF
FDR
INSTANTBACKUP
FOR ABR**

With FDR InstantBackup for ABR, all of the same information must be recorded on the online volume, to document the backup. However, these changes can't wait until the backup of the offline point-in-time image to tape is complete, since changes may have occurred to the online volume in the meantime (such as updating a data set again, deleting or allocating data sets).

To solve this, when ABR is invoked to capture the point-in-time image, it will update the online volume immediately. The ABR generation and cycle of the next backup is determined and all of the steps shown above are taken, except that the backup is cataloged to the online volume itself (something that would never occur in a normal ABR backup) so this tells ABR that the backup is a point-in-time image. This functions as a true backup, since you can do restores from it.

Later, ABR will process the point-in-time image and create a normal ABR backup from the offline image, exactly like it would do the backup of the online volume if InstantBackup were not involved. It will recatalog the backup data set to the tape just created. At this point, the backup looks exactly like a normal ABR backup of the online volume.

Even before the backup is moved to tape, ABR can restore the entire volume from the point-in-time image.

**EXECUTING FDR
INSTANTBACKUP
FOR ABR**

To use FDR InstantBackup for ABR, you must divide your existing ABR volume backup jobs into two steps which may be executed in one job or in two separate jobs:

- 1) the first step will capture all the point-in-time images of the selected online volumes using the ABR statement SPLIT. This will normally take only a few seconds per volume selected. We recommend that updates to the online volumes be suspended as far as possible during this step. As soon as it finishes, updates to the online volumes can begin again, since the point-in-time captured image of the volume exists. To insure that the snapped volume is not accidentally brought online, ABR will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail
- 2) the second step will create the actual backup tapes, reading the captured point-in-time images. It will be identical to the original ABR step, except that ABR operand BCV= must be added to invoke this support. This ABR step can also include backups of volumes which do not have captured images; these volumes will be backed up normally (with a warning message). It also optionally releases the captured images when it is done with them by re-establishing the BCV to its online volume)

When you split the ABR step into two, each step should contain the same JCL and ABR control statements as the original, with a few changes. It is vital that all ABR operands be identical in both steps, especially the TYPE= operand, except as noted below (actually, the second step will fail if the TYPE= operand is different). The examples later in this section illustrate the changes required.

25.06 CONTINUED

**EXECUTING FDR
INSTANTBACKUP
FOR ABR
(continued)**

In the first step:

- replace the operation name DUMP on the first ABR control statement with SPLIT. The ABR SPLIT operation will automatically use "differential split" on EMC microcode level 5x65 and "instant split" on microcode level 5x66 and above; both of these improve the performance of the SPLIT.
- change all TAPEX and TAPEXX DD statements to DUMMY. **However, if the original TAPEX/xx DD statements contained RETPD= or EXPDT= operands, include the same operands on the DUMMY DD statements so that ABR can record the actual expiration of the backups. Similarly, if you coded the ABR operand RETPD= on the DUMP statement, include it on the SPLIT statement as well.**
- Innovation recommends the use of DSNENQ=USE and ENQ=RESERVE so that ABR can properly protect the volume and data sets during the few moments necessary to create the point-in-time image. However, unlike normal ABR operations, where ABR will still backup a data set or volume even if DSNENQ= finds active data sets, SPLIT will not create the point-in-time if it finds data sets in use, which will also cause the second step to fail trying to backup that image. You can override this by specifying the ENQERR=NO operand on the DUMP statement, or by adding statements such as:

```
SELECT DSN=dsname , DSNENQ=NONE
```

for data sets which you expect to be active at the time of the SNAP/SPLIT; this is supported even for SPLIT TYPE=FDR

- If the backup includes other volumes in other disk subsystems, or volumes for which you simply don't want to do a point-in-time backup, they should be omitted from this step, even though they can be included in the second step

In the second step:

- add to the DUMP statement the operands

```
BCV=USE          (read point-in-time image) or
BCV=(USE,RET)    (read point-in-time image/resync BCV when done):
```

BCV=USE is used when you intend to pair the BCV with a different online volume after the backup is complete; BCV=(USE,RET) is used when the BCV is permanently assigned to the online volume.:

- DSNENQ= and ENQ= will be ignored by ABR when dumping point-in-time volume images, since these would prevent access to the online data sets and prevent changes to the online volume's VTOC, not the offline point-in-time image. However, they will be honored if you backup other non-point-in-time volumes in the same step

Since the target volumes can hold only one point-in-time image at a time, you must be sure that another ESTABLISH or RE-ESTABLISH is not done on the split BCV before the second step moves the backup to tape. If you don't, the backup of the previous point-in-time image will be lost.

If you have SPLIT a number of volumes, and then executed a DUMP with BCV=, but that DUMP step failed for some reason before processing all of the volumes, you can just resubmit the job. ABR will bypass the backup of any volume for which SPLIT was done and the backup has already been moved to tape. However, if you have included in the DUMP step other volumes which are to be backed up normally (e.g., volumes not capable of Timefinder), they will be backed up again by the restarted backup step.

CONTINUED . . .

25.06 CONTINUED

**MULTIPLE
SYSTEMS**

Although the SPLIT step and the DUMP step can execute on different MVS systems, **FDR InstantBackup for ABR requires that both systems must be able to access both the source and target volumes and a common ABR catalog.** During the SPLIT, ABR updates the online volume, the offline BCV copy and the ABR catalog. During the DUMP, ABR must read the online volume, and update the offline BCV copy and the ABR catalog.

It is not possible to SPLIT a BCV from one system, and then do the ABR backup on another system (either another LPAR or a remote system) unless that system can address both disk volumes and update the same ABR catalog. However, if you only need full-volume backups (no incremental backups), it is possible to do this outside of ABR with PGM=FDR backups; contact Innovation for guidance.

FDR InstantBackup for ABR does not currently support remote copies of volume using SRDF; the BCV must be in the same Symmetrix subsystem as its online volume.

25.07 EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK

These are examples of FDR InstantBackup usage with TimeFinder. The examples include steps which execute EMCTF, the EMC-provided TimeFinder utility, but you should consult current EMCTF documentation from EMC to verify the JCL and control statement requirements for that product.

A STEPLIB DD may be required in FDR InstantBackup steps if you have installed the product in a library other than your standard FDR program library.

**FDR FULL
VOLUME
BACKUP**

The BCV at address 01FA has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous one-time ESTABLISH has been issued to establish that pairing. The step SPLIT will split the BCV from its online volume and wait for the split to complete. Step BACKUP will backup the BCV copy of volume PROD01; when the dump finishes, FDR will re-synchronize the BCV with the online volume.

```
//SPLIT EXEC PGM=EMCTF
//SYSOUT DD SYSOUT=*
//SYSIN DD *
    SPLIT 1,01FA,WAIT
//BACKUP EXEC PGM=FDR,REGION=0M,COND=(0,NE,SPLIT)
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR.USE.UNIT01FA,UNIT=SYSALLDA,
//          VOL=SER=PROD01,DISP=OLD
//TAPE1 DD DSN=BACKUP.VPROD01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSIN DD *
    DUMP TYPE=FDR,BCV=(USE,RET)
```

The output from the backup step will look something like this (the FDR219 message will confirm that the offline BCV was actually backed up):

```
FDR303 CARD IMAGE -- DUMP TYPE=FDR
FDR219 FDR VOL=PROD01 IS BEING DUMPED FROM OFFLINE UNIT=01FA
FDR007 STARTING TIME OF ACTUAL DUMP -- 14.20.27 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR007 ENDING TIME OF ACTUAL DUMP -- 14.20.38 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR122 BYTES DSK TRK T BLKS RESTART STIMERS ERRS ACT DSK LOW HGH DEXCP NUMDS
FDR122N 0022053561 000618 000491 0000000 0000000 000 000618 000 000 00100 00010
FDR002 FDR DUMP SUCCESSFULLY COMPLETED FROM VOL=PROD01
FDR999 FDR SUCCESSFULLY COMPLETED
```

The backup created by this example is a normal FDR full-volume backup. There are no special considerations for restoring the entire disk volume or individual data sets from it. For example:

```
//RESTORE EXEC PGM=FDR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPE1 DD DSN=BACKUP.VPROD01(0),DISP=OLD
//DISK1 DD UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
//SYSIN DD *
    RESTORE TYPE=FDR
```

25.07 CONTINUED

**FDRDSF DATA SET
BACKUP**

The BCVs at addresses 01F2 and 01F3 are part of a pool of BCVs. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step ESTAB may be executed separately before the backup is required, giving the Symmetrix time to synchronize the BCV with the online volume; the step will terminate when the synchronization is complete. At that point updates to the online volumes can be quiesced and the rest of the steps executed. The step SPLIT will split the BCVs from their online volumes and wait for the split to complete; when it is complete, updates to the online volumes can resume. Step BACKUP will backup the BCV copies of selected data sets on volumes CICS01 and CICS02, processing the two backups in parallel.

```
//ESTAB    EXEC    PGM=EMCTF
//SYSOUT   DD      SYSOUT=*
//SYSIN    DD      *
      ESTABLISH 1,01F2,01D3,WAIT
      ESTABLISH 1,01F3,01D7,WAIT
//SPLIT    EXEC    PGM=EMCTF
//SYSOUT   DD      SYSOUT=*
//SYSIN    DD      *
      SPLIT 1,01F2-01F3,WAIT
//BACKUP   EXEC    PGM=FDRDSF,REGION=2M,COND=(0,NE,SPLIT)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1    DD      DSN=FDR.USE.UNIT01F2,UNIT=SYSALLDA,
//          VOL=SER=CICS01,DISP=OLD
//TAPE1    DD      DSN=BACKUP.VCICS01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN1 DD      SYSOUT=*
//DISK2    DD      DSN=FDR.USE.UNIT01F3,UNIT=SYSALLDA,
//          VOL=SER=CICS02,DISP=OLD
//TAPE2    DD      DSN=BACKUP.VCICS02(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN2 DD      SYSOUT=*
//SYSIN    DD      *
      DUMP TYPE=DSF,ATTACH
      SELECT DSN=CICS**
```

The backups created by this example are normal FDRDSF data set backups. There are no special considerations for restoring data sets from them. For example:

```
//RESTORE  EXEC    PGM=FDRDSF,REGION=2M
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//TAPE1    DD      DSN=BACKUP.VCICS01(0),DISP=OLD
//SYSIN    DD      *
      RESTORE TYPE=DSF,RECAT,VRECAT
      SELECT DSN=CICS3.LOADLIB
```

25.07 CONTINUED

**FDRCOPY FROM A
POINT-IN-TIME IMAGE**

An SMS storage group consists of 2 volumes, SMS001 and SMS002, which have permanently assigned BCVs at addresses 01F4 and 01FA. You want to make a copy of data sets on those volumes but you can only quiesce updates for a few moments. Step DEFINE will define the catalog entries which will allow FDR InstantBackup to access the BCVs (you can omit the DEFINE step if you have permanently cataloged the FDR.USE.UNITuuuu data sets). The step SPLIT will split the BCVs from their online volumes and wait for the split to complete. Step COPY will copy the selected data sets from the BCVs, renaming the data sets during the copy; the data sets will be copied to online volumes chosen by SMS as described in [Section 21](#). These copies might be used for parallel batch processing. After the copy is complete, FDRCOPY will re-synchronize the BCVs with the online volumes.

```
//DEFINE      EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
      DELETE FDR.USE.UNIT01F4 NOSCRATCH
      DELETE FDR.USE.UNIT01FA NOSCRATCH
      SET MAXCC=0
      DEFINE NONVSAM(NAME(FDR.USE.UNIT01F4) -
        DEVICETYPE(3390) VOLUME(SMS001))
      DEFINE NONVSAM(NAME(FDR.USE.UNIT01FA) -
        DEVICETYPE(3390) VOLUME(SMS002))
//SPLIT      EXEC PGM=EMCTF
//SYSOUT DD   SYSOUT=*
//SYSIN DD    *
      SPLIT 1,01F4,WAIT
      SPLIT 1,01FA,WAIT
//COPY      EXEC PGM=FDRCOPY,REGION=0M,COND=(0,NE,SPLIT)
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//DISK1 DD    DSN=FDR.USE.UNIT01F4,DISP=OLD
//DISK2 DD    DSN=FDR.USE.UNIT01FA,DISP=OLD
      COPY TYPE=DSF,(BCV=USE,RET)
      SELECT DSN=PROD.PAYROLL.***,NEWI=...+COPY,STORCLAS=BATCH
```


25.07 CONTINUED

**FDRCOPY
BETWEEN
SYMMETRIX
VOLUMES**

FDRCOPY can copy or move data sets to new volumes, and can rename and/or recatalog the output data sets. When FDR InstantBackup is installed, FDRCOPY is enhanced to automatically recognize that the input and output volumes for a given data set are in the same Symmetrix and automatically invokes an internal Symmetrix copy facility to do a fast copy. When the volumes for a given data set are not in the same Symmetrix, normal read/write is used. This function is automatic, requiring no special operands, so the samples below are normal FDRCOPY jobs. If the copy facility was used for a given data set, its FDR311 message will contain the text "INSTANT".

This example copies data sets to new names, inserting a second-level index of COPY.

```
//COPYFILE EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
COPY TYPE=DSF,MAXCARDS=1000
SELECT CATDSN=ABC.**,NEWI=+.COPY,
        NVOL=(TST001,TST002)
```

This example moves data sets to new volumes, recataloging them to their new location.

```
//MOVEFILE EXEC PGM=FDRCOPY,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
MOVE TYPE=DSF
SELECT CATDSN=MY.DATASET,NVOL=3390T1
SELECT CATDSN=*.DSNDBC.*.*.10001.A*,NVOL=3390T1
```

FASTCPK

Normally FASTCPK must read every data track being moved, and write it to its new location. When EMCCOPY=YES is specified and the volume being COMPAKed is in a EMC Symmetrix with the proper microcode, an EMC facility is used to copy data tracks from one location to another internally. The data is not sent to or from the host. This usually results in a reduction in elapsed time. If the volume is not in a Symmetrix, or the Symmetrix does not have the proper microcode level, normal read/write is done. This example will release space from all sequential, partitioned and VSAM data sets that have a secondary allocation quantity specified.

```
//FASTCPK EXEC PGM=FDRCPK,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSMAP DD SYSOUT=*
//SYSSUMM DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
COMPAKT TYPE=FASTCPK,VOL=EMC001,EMCCOPY=YES,
        CPKFREEEX=20,LOG=YES,NOSECOND=NORLSE,
        PSRLSE=ALL,PORLSE=ALL,VSRLSE=ALL
```

Note: if you simulate FASTCPK on a Symmetrix volume, the estimated elapsed time displayed by the simulation is calculated based on standard disk I/O times, not the reduced times provided by EMCCOPY=YES. Your actual elapsed time with EMCCOPY=YES will be significantly better than the simulation estimate.

25.08 EXAMPLES FOR ABR

ABR FULL VOLUME BACKUP BCVs have been permanently assigned to the online payroll volumes PAYxxx; a previous one-time ESTABLISH has been issued to establish those pairings.

Step FULL1 will split each BCV from its online volume to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
              SPLIT TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
              MOUNT VOLG=PAY*
```

The SYSPRINT output of step FULL1 will look something like:

```
FDR304  ABR SPLIT   REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305    TO OFFLINE UNIT=07C8 ,DSNAME=FDRABR.VPAY001.C1000300 ,FILE=0000,VOL=SER=PAY001
FDR306  SPLIT      SUCCESSFULLY COMPLETED
FDR304  ABR SPLIT   REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305    TO OFFLINE UNIT=07C9 ,DSNAME=FDRABR.VPAY002.C1000300 ,FILE=0000,VOL=SER=PAY002
FDR306  SPLIT      SUCCESSFULLY COMPLETED
FDR304  ABR SPLIT   REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305    TO OFFLINE UNIT=07CA ,DSNAME=FDRABR.VPAY003.C1000300 ,FILE=0000,VOL=SER=PAY003
FDR306  SPLIT      SUCCESSFULLY COMPLETED
```

The FDR304 message identifies this as a SPLIT, and the FDR305 message identifies the BCV target unit address (ABR queries the Symmetrix to get the address of the BCV currently associated with each online volume). It also shows the ABR backup data set name created, and that it was cataloged to the input volume.

The SYSPRINx output will also contain a message to indicate that a SPLIT was done:

```
FDR218    FDRBCV UNIT=07C0 IS SPLIT FROM BCV UNIT=07C8
```

25.08 CONTINUED

Step FULL2 will do the actual backups. BCV=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the Symmetrix hardware to get the addresses of the BCVs; there is no need to specify them here.

```
//FULL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11     DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2      DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22     DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN      DD     *
              DUMP TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,BCV=(USE,RET)
              MOUNT VOL=PAY*
```

The SYSPRINT output for step FULL2 will be similar to:

```
FDR219 FDR VOL=PAY001 IS BEING DUMPED FROM OFFLINE UNIT=07C8
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY001.C1000300,FILE=0005,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY001.C2000300,FILE=0001,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RE-CONNECTED TO BCVUNIT=07C8
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY002 IS BEING DUMPED FROM OFFLINE UNIT=07C9
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY002.C1000300,FILE=0012,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY002.C2000300,FILE=0001,VOL=SER=B15600
FDR218 FDRINST UNIT=07C2 WAS RE-CONNECTED TO BCVUNIT=07C9
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY003 IS BEING DUMPED FROM OFFLINE UNIT=07CA
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY003.C1000300,FILE=0006,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY003.C2000300,FILE=0002,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RE-CONNECTED TO BCVUNIT=07CA
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR224 FDR VOL=PAY004 IS BEING DUMPED FROM ONLINE UNIT BECAUSE SNAP/SPLIT WAS NOT ATTEMPTED
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY004,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY004.C1000300,FILE=0013,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY004.C2000300,FILE=0002,VOL=SER=B15600
FDR306 DUMP SUCCESSFULLY COMPLETED
```

FDR219 messages indicate that a previously created point-in-time image was found for the volume being dumped. FDR305 messages show the actual location of the backup data sets. FDR218 messages indicate that the RET option caused ABR to RE-ESTABLISH the BCV as an active mirror of its online volume, ready to do the next backup.

25.08 CONTINUED

In this example, volume PAY004 has no BCV (a BCV was never ESTABLISHED for this volume). The FDR224 message indicates that there is no point-in-time image of the volume so that a normal ABR backup was done. If you intended to do point-in-time backups all of the volumes in this backup, you must do a EMCTF ESTABLISH for PAY004 and insure that the FULL1 step includes it; otherwise the backup from PAY004 will not be taken at the same point-in-time as the other volumes. If you have no need to do point-in-time backup on PAY004, you can ignore the message. So, you can include in this ABR step volumes which have no point-in-time requirements; they will be backed up normally.

Note that if you attempted to SPLIT a volume but the SPLIT failed for some reason (perhaps because the volume is not in a Symmetrix), ABR will set a flag in the ABR Model DSCB on the online volume. The DUMP step will fail the backup of that volume with message FDR211 and a non-zero return code, assuming that you intended it to be a point-in-time backup.

**ABR
INCREMENTAL
BACKUP**

The BCVs at addresses 01F2 and 01F3 are part of a pool of BCVs. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step ESTAB may be executed separately before the backup is required, giving the Symmetrix time to synchronize the BCV with the online volume; the step will terminate when the synchronization is complete. At that point updates to the online volumes can be quiesced and the rest of the steps executed.

```
//ESTAB EXEC PGM=EMCTF
//SYSOUT DD SYSOUT=*
//SYSIN DD *
ESTABLISH 1,01F2,01D3,WAIT
ESTABLISH 1,01F3,01D7,WAIT
```

Step INCR1 will split the BCV from its online volume to create a point-in-time image and mark the backup as complete. It will create a new ABR backup cycle and update the online volume with information about the new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that the RETPD= parameter must be specified the same as in the INCR1 step which follows.

```
//INCR1 EXEC PGM=FDRABR,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSPRIN1 DD SYSOUT=*
//SYSPRIN2 DD SYSOUT=*
//TAPE1 DD DUMMY
//SYSIN DD *
SPLIT TYPE=ABR,DSNENQ=USE,RETPD=14,ENQERR=NO
MOUNT VOL=CICS01
MOUNT VOL=CICS02
```

Step INCR2 will do the actual backups. BCV=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the Symmetrix hardware to get the addresses of the BCVs; there is no need to specify them here.

```
//INCR2 EXEC PGM=FDRABR,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSPRIN1 DD SYSOUT=*
//SYSPRIN2 DD SYSOUT=*
//TAPE1 DD DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP)
//SYSIN DD *
DUMP TYPE=ABR,DSNENQ=USE,DATA=USED,RETPD=14,ENQERR=NO,BCV=(USE,RET)
MOUNT VOL=CICS01
MOUNT VOL=CICS02
```

For examples of messages and other notes, see the preceding example, "ABR FULL VOLUME BACKUP"

CONTINUED . . .

25.08 CONTINUED

ABR RESTORE

If it is necessary to do a full-volume restore of one of the volumes involved in the point-in-time backup before the backup is moved to tape, ABR will do it automatically. No special parameters or JCL is required. ABR will restore from the point-in-time image or from tape automatically.

For example, this will restore the four volumes in the previous full-volume backup example:

```
//REST      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT  DD     SYSOUT=*
//SYSIN     DD     *
RESTORE TYPE=FDR,DYNTAPE,CONFMESS=NO,ONLINE,CPYVOLID=YES
SELECT VOL=PAY001
SELECT VOL=PAY002
SELECT VOL=PAY003
SELECT VOL=PAY004
```

Volumes PAY001-003 will be restored from the point-in-time images on the BCVs unless they have been moved to tape. Volume PAY004, since it has no point-in-time image, will always be restored from tape.

However, there are some considerations:

- the point-in-time volume image is always a full-volume image, even if a incremental backup was done. If ABR discovers that the first ABR backup cycle to be restored is a point-in-time image, it simply uses the appropriate hardware features to quickly copy the entire volume back to the target online volume and does not need to read any further backups from tape. Since full-volume restores typically start with the most recent backup of the volume, this provides the fastest recovery.
- but if the most recent ABR backup cycle is on tape, but another lower-numbered cycle is a point-in-time image, ABR will bypass restoring from that image (with a message); any data sets whose most recent backup was on that image may be corrupted or back-level. This could occur if you did a SPLIT to create a cycle but never did a DUMP to move it to tape, and followed that with a regular ABR backup direct to tape. This is not normal and should not occur if your FDR InstantBackup procedures are setup correctly.
- If you are restoring to a volume other than the original volume (NVOL= on the SELECT statement) and the target volume is not in the same disk subsystem as the point-in-time image, ABR will do a normal disk-to-disk copy to copy all tracks in the image to that volume.
- If ABR attempts to restore from a point-in-time image (ABR backup data set cataloged to the original volume), but that image no longer exists, ABR will bypass that backup. If it was an incremental backup, it will restore the rest of the cycles in that generation. If it was a full-volume backup (cycle 00), the restore will fail (you can uncatalog the backup or specify GEN= and CYCLE= to restore from an earlier generation). This can occur if multiple SPLIT operations were done without an intervening DUMP to move the backup to tape. It can also occur at a disaster site if the restored ABR catalog is not consistent with the tapes that were brought off-site (some backup entries still indicate the point-in-time backup).

26.01 FDR INSTANTBACKUP FOR SNAPSHOT

**FDR
INSTANTBACKUP
AND IBM
SNAPSHOT**

FDR InstantBackup is an additional cost FDR facility for non-disruptive backup. It works in concert with Snapshot, a feature of the IBM 9393 RVA (RAMAC Virtual Array) and StorageTek Iceberg and SVA (Shared Virtual Array) which allows for the "instant" creation of exact duplicates of existing DASD volumes.

FDR InstantBackup enhances:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY data set copy
- FASTCPK volume reorganization
- FDRABR volume backup (full-volume and incremental backup)

It provides non-disruptive backup of offline MVS DASD volumes, both SMS and non-SMS, using technology developed by Innovation Data Processing.

SNAPSHOT

Traditional backups may be disruptive of normal operations. It obviously takes time to backup a data set or volume. If data sets are being updated by some application while they are being backed up, the backup may not be valid and the data sets may not be usable when they are restored. The usual response to this problem is to quiesce all updates to the data during the backup, disrupting normal operation until the backup is complete.

When used with the Snapshot feature of the RVA, Iceberg, or SVA, FDR InstantBackup allows you to instantly freeze a copy of an online disk volume and easily backup or copy data from that frozen copy as if it was coming from the original volume. Updates to the data need to be quiesced only for the few moments necessary to create the frozen copy.

**FDR
INSTANTBACKUP
ENHANCES FDR
COMPONENTS**

FDR InstantBackup enhances FDR, FDRDSF, and FDRCOPY to read from an offline copy of an online volume which was created as a point-in-time frozen image using Snapshot. This creates a backup or copy of the volumes or data sets that looks exactly like a normal backup or copy, except that the data is frozen in time.

FDR InstantBackup enhances FDRCOPY to use Snapshot. When copying or moving data sets with FDRCOPY, if the input volume and the output volume for a given data set are both within the same disk subsystem, FDRCOPY will issue direct Snapshot commands to copy that data set, resulting in almost instantaneous copies. If the input and output volumes for a data set are not in the same disk subsystem or the Snapshot feature is not installed, normal reads and writes are done to copy the data. This feature is automatic; no special options are required to invoke it.

If you are licensed for COMPAKTOR, FDR InstantBackup also enhances FASTCPK to use the Snapshot facility to move tracks internally without sending the track contents to the host and back again (READ and WRITE). If you specify SNAPSHOT=YES on the COMPAKT statement, FASTCPK will use Snapshot to move tracks on the volume being COMPAKed, resulting in a significant savings in elapsed time.

If you are licensed for FDRABR, FDR InstantBackup enhances ABR Volume Backups, both full-volume (TYPE=FDR) and incremental (TYPE=ABR/DSF/AUTO) to use Snapshot, by creating "instant" frozen copies of the volumes to be backed up, and then moving that captured data to tape.

FDR InstantBackup also supports HSDM (High Speed Data Mover), which is a hardware option on StorageTek SVAs and on enhanced IBM RVAs. [See Section 80.33](#) for more details.

CONTINUED . . .

26.02 SNAPSHOT OPERATION

The IBM RVA and StorageTek Iceberg and SVA all use a method of storing data which is currently unique among mainframe storage vendors. It does use RAID storage to provide for reliability and recoverability but adds an additional function.

Like other mainframe disks and RAID systems, these subsystems emulate traditional IBM disk devices, namely 3390 and 3380. Unlike other RAID systems, there are no fixed locations for the data tracks belonging to the emulated disks. In fact, any updated data track is always written to a new location within the subsystem, and data tracks belonging to various 3390/3380 volumes will actually be intermingled on the physical disks within the subsystem. Data tracks are also compressed.

The subsystem keeps track of this with a table which maps all of the defined 3390 and 3380 disk volumes (called "functional disks") to the current location of each data track on each volume. For tracks which have never been written, or which are not currently assigned to any data set on the functional volume, the pointer is zero and they take up no space in the subsystem.

Please consult IBM or StorageTek for more information on the operation and features of the IBM RVA or StorageTek Iceberg or SVA.

SNAPSHOT SNAPSHOT is an additional cost feature of these subsystems, which allows for the instant duplication of data sets or entire disk volumes, from one functional volume to another functional volume in the same subsystem.

When Snapshot is invoked, the pointers to the data tracks of the source data are simply copied into the table for the target data. The original data tracks are now addressable by two addresses. Either address will actually access the same data stored on the internal disks.

The "snapped" copy of a track takes up no extra space in the subsystem, unless some program updates the track either by its original (source) address or its new (target) address. In this case, the new data track is written elsewhere in the subsystem and the pointer for the address used to access it is updated, but the pointer for the other copy still points to the original data and is not updated. So, once data is copied with Snapshot, updates can take place against the original data but the copied data is not changed, giving you a true "instant" image of the copied data.

Here are two examples to help you understand how Snapshot works:

- 1 You request that Snapshot copy an entire disk volume at address 142 to another volume in the same subsystem at address 152. The subsystem copies all of the disk track pointers from disk 142 to the track table for disk 152; they point to the same data. A read for data track 10 on either volume will read the same data from the internal disks or cache.
- 2 You request that Snapshot copy a data set (A) from some volume to another data set (B) in the same subsystem. The subsystem copies all of the data track pointers for the tracks belonging to data set A on top for the pointers for data set B. Note that A and B may be on the same volume or on different volumes. The two data sets are probably at different locations on their respective volumes. If either data set A or B is opened for input, the same data will be read by the subsystem.

IBM and StorageTek provide a utility program IXFP for the management of these subsystems. If you have the hardware Snapshot feature on your subsystem you will also have the Snapshot feature of IXFP in order to issue Snapshot commands. However, FDR InstantBackup includes all the necessary support for Snapshot; IXFP must be installed but it is not directly used by FDR InstantBackup.

Note: Snapshot cannot be used if the target volume is currently part of a PPRC (Peer-to-Peer Remote Copy) pair; this is a microcode restriction. FASTCPK and FDRCOPY will use standard I/O if the target volume is part of a PPRC pair, causing the operation to take longer than expected; the FDRSNAP and FDRABR SNAP commands will fail in this case.

CONTINUED . . .

26.02 CONTINUED . . .

**SNAPSHOT
WITHOUT FDR
INSTANT BACKUP**

Although Snapshot can instantly create a point-in-time image of a production disk volume, traditional backup products cannot directly use the snapped image without additional considerations.

For example, remember that a full-volume snap will duplicate pointers to all of the original volume's data tracks, including the label track containing the volume serial. Since MVS will not let you put two volumes with the same volume serial online, the new volume is forced offline. A conventional backup product will require that the duplicated volume be brought online, so the volume serial of the snap target volume must be changed.

IXFP Snapshot support provides a facility for re-labeling the target volume during the full-volume snap so that it can be put online, but re-labeling a volume invokes additional requirements for SMS managed volumes, VSAM clusters and cataloged data sets. The data sets on the re-labeled volume will appear to be uncataloged, which may cause problems during backup and restore unless they are renamed and re-cataloged.

A backup of a re-labeled volume will appear to be a backup from the new volume serial, not the original volume serial, so additional procedures are required to document this correspondence (e.g., backup of volume B is really a backup of volume A). Restore procedures must be modified to restore the backups to the correct serial (e.g., restore volume B back to volume A and re-label it as A).

If you use the data set copy feature of Snapshot, the copied data sets will usually have new names. Although this could be used to backup those data sets with a standard backup utility, the installation must now use special restore procedures to restore each data set from its temporary name on the backup to its true name on the online volume. For large numbers of data sets, this could be onerous.

26.03 THE FDR INSTANTBACKUP SOLUTION

FDR InstantBackup allows you to create a duplicate volume image with Snapshot **without changing its volume serial or bringing it online**. FDR full volume backup, FDRDSF data set backup, FDRCOPY data set copy and FDRABR volume backups use FDR InstantBackup technology to read the duplicated volume even while it remains offline. A non-disruptive backup or copy can start as soon as the snap is done. Updates to the original volumes can continue while you create the point-in-time backup or copy of the data.

For FDRABR, FDR InstantBackup makes volume backups a 2-step process. The first step executes a SNAP statement under FDRABR, to copy each selected volume to an offline disk in the same RVA/Iceberg/SVA at an address that you provide. This takes only seconds per volume. The snapped copy is actually a valid backup at that point-in-time; you can do a restore from it (using Snapshot) if recovery is required. The second step executes a normal FDRABR full-volume or incremental backup of the same volumes with a parameter which tells ABR to read the offline point-in-time image and move it to a backup on tape. There is no need to identify the snapped device to the ABR backup step; it knows the address of the most recent copy of each online volume.

FDR InstantBackup also makes FDR or FDRDSF backups and FDRCOPY copies a 2-step process. The first step executes a SNAP statement under program FDRSNAP, to copy each selected volume to an offline disk in the same subsystem at an address that you provide. When that step ends, the data has been copied and you can resume normal operations on the original volume. Then you execute a FDR or FDRDSF backup or a FDRCOPY copy, with a parameter which tells FDR to read the snapped copy instead of the original volume. It then will read the offline copy in exactly the same way that it would read the online disk. FDR InstantBackup produces a backup that appears to be a conventional backup of the original volume serial. FDR InstantBackup blends its unique off-line, high-speed non-disruptive backup together with a traditional restore complete with FDR's powerful logical file capability.

There are no special considerations for restores from a backup created by FDR InstantBackup. You can restore entire volumes or individual data sets from those backups. The target volumes for restore will be normal online volumes. Since the backups created by FDR InstantBackup appear to be backups of the original volume, there are no special volume serial or data set name concerns.

FDRCOPY data set copy and move can be used with FDR InstantBackup in two ways:

- when FDRCOPY detects that the input and output volume associated with a given data set are in the same subsystem with the Snapshot feature installed, it automatically invokes the Snapshot hardware to copy the tracks, resulting in an almost instantaneous copy. Normal read/write is used if the volumes are not both in the same subsystem or if the feature is not installed.
- Alternately, you can use Snapshot to freeze an image of the volumes containing the desired data sets, and then use FDRCOPY to create copies of those data sets on other online volumes. This is useful if the target volumes are not in the same subsystem as the source volumes, allowing you to capture an instant image of the source data sets and copy them at your leisure.

CONTINUED . . .

26.04 SETTING UP SNAPSHOT

**SNAPSHOT
CONFIGURATION**

Once you have installed the proper IBM or StorageTek features and IXFP software to support Snapshot, there are no other requirements for the use of Snapshot.

However, you will want to reserve a set of device addresses (functional volumes) in each subsystem for use as Snapshot target volumes for FDR InstantBackup. Since a full-volume snap destroys any data already on the target volume, these volumes cannot be used for any other type of storage.

The good news is that these Snapshot target volumes will not occupy any physical disk space in the subsystem. When not in use, they occupy no space at all. When used as a Snapshot target, they still occupy no additional space, since all the data track pointers point to the original tracks. Only if a significant number of data tracks on the original volume are updated before FDR InstantBackup completes will these volumes have any impact on the NCL (Net Capacity Load, the total data stored in the subsystem) but this is short-lived.

The device addresses assigned to the Snapshot targets must be defined as normal 3380 or 3390 disks in the disk subsystem in your I/O configuration (IOCP or HCD).

When these disks are in use with FDR InstantBackup, using the SNAP TYPE=FDR statement, FDR will modify the volume label of the snapped volume to insure that it cannot be brought online. To avoid the error messages during IPL due to these modified labels, you may want to mark the Snapshot target addresses as OFFLINE in your I/O configuration, so that MVS will not attempt to access them at IPL time. FDR InstantBackup will be able access them even though they are offline.

For example, in HCD, the MVS options for a Snapshot target device will look similar to:

View Device Parameter / Feature DefinitionRow 1 of 6

Command ==> Scroll ==> PAGE

Configuration ID : CURP1099Common Configuration data

Device number : 01E0Device type : 3390

Generic / VM device type : 3390

Parameter/	Value	Req.	Description
Feature			
OFFLINE	Yes		Device considered online or offline at IPL
DYNAMIC	Yes		Device supports dynamic configuration
LOCANY	Yes		UCB can reside in 31 bit storage
ALTCTRL	No		Separate physical control unit path
SHARED	Yes		Device shared with other systems
SHAREDUP	No		Shared when system physically partitioned

In IOCP/MVSCP, specify OFFLINE=YES on the IODEVICE macros for the Snapshot target devices. The following is an example of a snapshot-capable subsystem attached via 2 ESCON channels (CHPIDs 40 and 41), with 128 online disk addresses (300-37F) and 64 Snapshot addresses (380-3BF):

```

CNTLUNIT CUNUMBR=008,PATH=(40),LINK=(C0),
        UNIT=3990,UNITADD=((00,192))
CNTLUNIT CUNUMBR=009,PATH=(41),LINK=(C4),
        UNIT=3990,UNITADD=((00,192))
IODEVICE ADDRESS=(300,128),CUNUMBR=(008,009), ** ONLINE VOLUMES
        UNIT=3390,UNITADD=00,FEATURE=SHARED
IODEVICE ADDRESS=(380,64),CUNUMBR=(008,009), ** SNAPSHOT TARGET VOLUMES
        UNIT=3390,UNITADD=80,FEATURE=SHARED,OFFLINE=YES

```

Note that some IOCP documentation may indicate that UNITADD=00 is required for ESCON; what it really means is that the value for the **first** IODEVICE definition for a given control unit must be 00. So when varying characteristics must be specified for some devices on a control unit (such as OFFLINE=YES), you can have multiple IODEVICE statements with varying values for UNITADD.

CONTINUED . . .

26.04 CONTINUED . . .

**SNAPSHOT WITH
FDR
INSTANTBACKUP**

At any time you can take a full-volume Snapshot of an online volume to one of the Snapshot target volumes you have reserved, using the SNAP statement of program FDRSNAP (for FDR, FDRDSF, or FDRCOPY) or FDRABR (for ABR volume backups). The duplication is almost instantaneous.

If you plan to backup multiple volumes and especially if multiple volumes contain related data which must be backed up at the same point-in-time, you should snap all of those volumes at one time.

As soon as the SNAP step ends, the data has been preserved and you can begin updating the original volumes while the backups are run.

The snapped copy will not add to the amount of data actually written on the disks within the subsystem (the NCL "Net Capacity Load") unless the source tracks are updated. Since this becomes increasingly likely as time goes on, FDR InstantBackup provides an option to release the storage associated with the snapped volumes as soon as the backup of each volume is complete.

You never need to bring a snapped copy online or re-label it or re-initialize it with ICKDSF, as long as it is used only for FDR InstantBackup.

**SETTING UP A
SNAPSHOT
ENVIRON-
MENT**

There are 2 ways to setup Snapshot use with FDR InstantBackup:

- 1) you can permanently assign a Snapshot target volume to every original volume for which you plan to use FDR InstantBackup. This will require more device addresses, but is much simpler to administer, since you always know which Snapshot target will be used for each original volume.
- 2) you can create a pool of Snapshot targets and use them with FDR InstantBackup as necessary. This requires fewer device addresses, but is more complex to administer. You will require operational procedures to insure that two FDR InstantBackup jobs don't try to use the same target volume for two different original volumes at the same time.

You can use a combination of the two, some online volumes with permanent targets (perhaps for volumes with frequent or time-critical backups) and some with targets assigned from a pool (perhaps for less frequent or non-time-critical backups).

In either case, the target volumes must match the size (number of cylinders) and format (3380 or 3390) of the original volumes to which they will be assigned. The target volume can be larger than the original volume (more cylinders).

26.05 FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY

To create the Snapshot copies of the volumes to be processed by FDR, FDRDSF or FDRCOPY, you must first execute program FDRSNAP with JCL similar to:

```
//SNAPVOLS EXEC PGM=FDRSNAP
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DUMMY
//SYSIN DD *
        SNAP TYPE=FDR
        MOUNT VOL=PROD01, SNAPUNIT=037D
        MOUNT VOL=PROD02, SNAPUNIT=1D97
```

In the SNAP step, one or more MOUNT statements must be provided to identify each online source volume to be copied (VOL=) and the target device (SNAPUNIT=) to which to snap it. There must be one MOUNT statement for each volume to be snapped, and the UNIT= value must be unique for each source volume. The online volume identified by VOL= and the target volume identified by SNAPUNIT= must be in the same subsystem and the target must be offline to MVS. To insure that the snapped volume is not accidentally brought online, FDRSNAP will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail.

WARNING: 4-digit device addresses must be used with SNAPUNIT= even if your operating system only supports 3-digit addresses for other functions; specify a leading zero if required.

Ideally, you should quiesce any updates to the source volumes during the snap, to insure that you can get a clean, restorable backup of every data set on the volumes. Once the FDRSNAP step ends (usually within a few seconds), the data has been copied and normal access to the source volumes can be resumed.

FDRSNAP does the same security checks as an FDR full-volume backup (see [Section 10.02](#)).

If you like, you can specify the DSNENQ= operand (to check if data sets are active when snapped) and/or the ENQ= operand (to ENQ or RESERVE the VTOC during the snap). See [Section 10.04](#) for details.

The actual backups of the offline point-in-time copies can now be taken with FDR InstantBackup using:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY input disks for COPY operations only.

You can also use FDR InstantBackup for FDR full-volume copies (COPY TYPE=FDR) but in most cases it would be more useful to copy the online volume directly.

FDRCOPY MOVE operations from a point-in-time copy are not supported by FDR InstantBackup, since it deletes the input data set. You should always move the live, online version of a data set.

FDR InstantBackup uses standard FDR JCL and control statements with one modification. For complete details on other JCL and control statement requirements of FDR, FDRDSF, and FDRCOPY, please see [Sections 10, 20, and 21](#) of the FDR manual, respectively. DISKx DD statements in FDR InstantBackup steps must specify the original online disk volume, just as if you were backing it up directly.

CONTINUED . . .

26.05 CONTINUED . . .

To invoke FDR InstantBackup you must add the SNAP= operand to the DUMP or COPY statement. SNAP= has two forms:

SNAP=USE - tells FDR InstantBackup to read the offline snapped copy of the online volume specified in JCL. FDR remembers the device address of the offline device most recently used as a target for a snap of each online volume unless an intervening IPL has occurred.

SNAP=(USE,REL) - same as SNAP=USE, but at the end of backing up each snapped volume, FDR will issue a request to the subsystem to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.

Note: FDR can determine the device address of the offline snapped copy **only** if FDRSNAP was used to create the snapped copy. If you use IXFP (SIBBATCH) or any other facility to invoke Snapshot, FDR will not be able to determine that address; contact Innovation for assistance if you must use IXFP.

Other considerations for FDR control statements with FDR InstantBackup:

- Since FDR InstantBackup is not reading the live data, it is never appropriate to use the ENQ= or DSNENQ= operands to enqueue on the VTOC of the Snapshot target or the data sets on the target. To do so would unnecessarily prevent access to the live data on the online volume.
- For FDRCOPY, specify DSNENQ=NONE. If you use CATDSN= to select data sets to be copied, you must make sure that all of the input volumes which will be selected from the catalog have been copied by a preceding FDRSNAP step.
- Do not use the FDRCOPY MOVE statement when the input is a Snapshot target. If you do so, the live data sets on the online volume will be deleted at the end of the move.

FDR INSTANTBACKUP OPERATION

- 1) if you do not have a permanently assigned Snapshot target volume for every original volume, select target volumes from the pool for use with this execution. You must establish procedures to insure that a given target is not used for another SNAP until the backup is complete.
- 2) When ready to backup or copy data, quiesce system and/or application update activity on the online volumes, if required, and execute FDRSNAP for each original-target pair. Once the FDRSNAP step ends, you can re-enable updates on the original volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with SNAP=USE or SNAP=(USE,REL) against each of the online volumes. The offline snapped copy will be read.

MULTIPLE SYSTEMS

FDR InstantBackup requires that both the original volume and its offline snapped copy be accessible from the MVS system where the backup is run. If you execute the SNAP on one system, but execute the backup on a second MVS system which does NOT have access to the online volume, FDR InstantBackup will fail. There is a circumvention for FDR and FDRDSF, but you must contact Innovation for assistance in implementing it.

Even if the volumes are accessible from both systems, the DUMP step will fail if it is run on a different system (another CPU or LPAR) from that where the SNAP was run. If this is required, contact Innovation for assistance.

26.06 FDR INSTANTBACKUP FOR ABR

When a normal ABR full-volume or incremental backup completes successfully, it is recorded by doing 2 things:

- ABR catalogs the backup data set it just created in the ABR catalog
- ABR updates the VTOC on the volume to record in the ABR Model DSCB the generation and cycle of the backup just created, and updates the current backup information for every selected data set

So, as soon as the backup data set is complete and available, all ABR records point to it and restores can be immediately requested.

**OPERATION OF
FDR
INSTANTBACKUP
FOR ABR**

With FDR InstantBackup for ABR, all of the same information must be recorded on the online volume, to document the backup. However, these changes can't wait until the backup of the offline point-in-time image to tape is complete, since changes may have occurred to the online volume in the meantime (such as updating a data set again, deleting or allocating data sets).

To solve this, when ABR is invoked to capture the point-in-time image, it will update the online volume immediately. The ABR generation and cycle of the next backup is determined and all of the steps shown above are taken, except that the backup is cataloged to the online volume itself (something that would never occur in a normal ABR backup) so this tells ABR that the backup is a point-in-time image. This functions as a true backup, since you can do restores from it.

Later, ABR will process the point-in-time image and create a normal ABR backup from the offline image, exactly like it would do the backup of the online volume if InstantBackup were not involved. It will recatalog the backup data set to the tape just created. At this point, the backup looks exactly like a normal ABR backup of the online volume.

Even before the backup is moved to tape, ABR can restore the entire volume from the point-in-time image.

**EXECUTING FDR
INSTANTBACKUP
FOR ABR**

To use FDR InstantBackup for ABR, you must divide your existing ABR volume backup jobs into two steps which may be executed in one job or in two separate jobs:

- 1) the first step will capture all the point-in-time images of the selected primary volumes using the ABR statement SNAP. This will take only a few seconds per volume selected. We recommend that updates to the online volumes be suspended as far as possible during this step. As soon as it finishes, updates to the online volumes can begin again, since the point-in-time captured image of the volume exists. To insure that the snapped volume is not accidentally brought online, **ABR** will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail
- 2) the second step will create the actual backup tapes, reading the captured point-in-time images. It will be identical to the original ABR step, except that ABR operand SNAP= must be added to invoke this support. This ABR step can also include backups of volumes which do not have captured images; these volumes will be backed up normally (with a warning message). It also optionally releases the captured images when it is done with them to reduce the NCL (Net Capacity Load).

When you split the ABR step into two, each step should contain the same JCL and ABR control statements as the original, with a few changes. It is vital that all ABR operands be identical in both steps, especially the TYPE= operand, except as noted below (actually, the second step will fail if the TYPE= operand is different). The examples later in this section illustrate the changes required.

CONTINUED . . .

26.06 CONTINUED . . .

EXECUTING FDR
INSTANTBACKUP
FOR ABR
(continued)

In the first step:

- replace the operation name DUMP on the first ABR control statement with SNAP
- change all TAPEX and TAPEXX DD statements to DUMMY. **However, if the original TAPEX/xx DD statements contained RETPD= or EXPDT= operands, include the same operands on the DUMMY DD statements so that ABR can record the actual expiration of the backups. Similarly, if you coded the ABR operand RETPD= on the DUMP statement, include it on the SNAP statement as well.**
- If the original ABR step used DISKxxxx DD statements or mount statements with VOLG= or STORGRP= to identify the volumes to be dumped, you must change the SNAP step to provide one MOUNT statement for each volume to be snapped, since the MOUNT statement provides the unit address of the target volume for each snap. DISKxxxx DD statements can be left in place; it will work with or without them.
- Innovation recommends the use of DSNENQ=USE and ENQ=RESERVE so that ABR can properly protect the volume and data sets during the few moments necessary to create the point-in-time image. However, unlike normal ABR operations, where ABR will still backup a data set or volume even if DSNENQ= finds active data sets, SNAP will not create the point-in-time if it finds data sets in use, which will also cause the second step to fail trying to backup that image. You can override this by specifying the ENQERR=NO operand on the DUMP statement, or by adding statements such as:

```
SELECT DSN=dsname,DSNENQ=NONE
```

for data sets which you expect to be active at the time of the SNAP; this is supported even for SNAP TYPE=FDR

- If the backup includes other volumes in other disk subsystems, or volumes for which you simply don't want to do a point-in-time backup, they should be omitted from this step, even though they can be included in the second step

For example,

```
//STEP1 EXEC PGM=FDRABR,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
        SNAP TYPE=FDR
        MOUNT VOL=ABC123,SNAPUNIT=03A4
        MOUNT VOL=XYZ321,SNAPUNIT=3B76
```

One or more MOUNT statements must be provided to identify each online source volume to be copied (VOL=) and the target device (SNAPUNIT=) to which to snap it. There must be one MOUNT statement for each volume to be snapped, and the UNIT= value must be unique for each source volume. The online volume identified by VOL= and the target volume identified by SNAPUNIT= must be in the same subsystem and the target must be offline to MVS. To insure that the snapped volume is not accidentally brought online, FDRSNAP will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail.

26.06 CONTINUED . . .

**EXECUTING FDR
INSTANTBACKUP
FOR ABR
(continued)**

In the second step:

- add to the DUMP statement the operands

SNAP=USE - tells ABR to read the offline snapped copy of the online volume specified in JCL. ABR remembers the device address of the offline device most recently used as a target for a snap of each online volume.

SNAP=(USE,REL) - same as SNAP=USE, but at the end of backing up each snapped volume, ABR will issue a request to the subsystem to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.

- DSNENQ= and ENQ= will be ignored by ABR when dumping point-in-time volume images, since these would prevent access to the online data sets and prevent changes to the online volume's VTOC, not the offline point-in-time image. However, they will be honored if you backup other non-point-in-time volumes in the same step
- If your disk subsystem also includes the HSDM (High Speed Data Mover) option, specify DCT=YES on the DUMP statement to reduce the elapsed time of the backup up to 60%.

Since a target volume can hold only one point-in-time image at a time, you must be sure that another SNAP is not done to the target volume before the second step moves the backup to tape. If you don't, the backup of the previous point-in-time image will be lost.

If you have snapped a number of volumes, and then executed a DUMP with SNAP=, but that DUMP step failed for some reason before processing all of the volumes, you can just resubmit the job. ABR will bypass the backup of any volume for which SNAP was done and the backup has already been moved to tape. However, if you have included in the DUMP step other volumes which are to be backed up normally (e.g., volumes not capable of Snapshot), they will be backed up again by the restarted backup step.

**MULTIPLE
SYSTEMS**

Although the SNAP step and the DUMP step can execute on different MVS systems, **FDR InstantBackup for ABR requires that both systems must be able to access both the source and target volumes and a common ABR catalog.** During the SNAP, ABR updates the online volume, the offline snapped copy and the ABR catalog. During the DUMP, ABR must read the online volume, and update the offline snapped copy and the ABR catalog.

It is not possible to SNAP a volume from one system, and then do the ABR backup on another system (either another LPAR or a remote system) unless that system can address both disk volumes and update the same ABR catalog. However, if you only need full-volume backups (no incremental backups), it is possible to do this outside of ABR with PGM=FDR backups; contact Innovation for guidance.

26.07 EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK

These are examples of FDR InstantBackup usage with Snapshot.

A STEPLIB DD may be required in FDR InstantBackup steps if you have installed the product in a library other than your standard FDR program library.

**FDR FULL
VOLUME
BACKUP**

The Snapshot target volume at address 01FA has been permanently assigned to online volume "PROD01". The step SNAP will snap PROD01 to 01FA. Step BACKUP will backup the snapped copy of volume PROD01. After the backup is complete, FDR will release all of the back-end (internal disk) storage associated with the snapped volume. DCT=YES requests that the HSDM (High Speed Data Mover) hardware option be used if it is installed on your disk subsystem; HSDM can reduce backup times up to 60%.

```
//SNAP      EXEC      PGM=FDRSNAP
//SYSPRINT DD          SYSOUT=*
//TAPE1     DD          DUMMY
//SYSIN     DD          *
      SNAP TYPE=FDR
      MOUNT VOL=PROD01, SNAPUNIT=01FA
//BACKUP    EXEC      PGM=FDR, REGION=0M, COND=(0, NE, SNAP)
//SYSPRINT DD          SYSOUT=*
//SYSUDUMP DD          SYSOUT=*
//DISK1     DD          UNIT=SYSALLDA, VOL=SER=PROD01, DISP=OLD
//TAPE1     DD          DSN=BACKUP.VPROD01(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSIN     DD          *
      DUMP TYPE=FDR, SNAP=(USE, REL), DCT=YES
```

The output from the backup step will look something like this (the FDR219 message will confirm that the offline snapped copy was actually backed up):

```
FDR303  CARD IMAGE -- DUMP TYPE=FDR
FDR219  FDR VOL=PROD01 IS BEING DUMPED FROM OFFLINE UNIT=01FA
FDR007  STARTING TIME OF ACTUAL DUMP -- 14.20.27 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR007  ENDING TIME OF ACTUAL DUMP -- 14.20.38 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR122  BYTES DSK TRK T BLKS RESTART STIMERS ERRS ACT DSK LOW HGH DEXCP NUMDS
FDR122N 0022053561 000618 000491 0000000 0000000 000 000618 000 000 00100 00010
FDR002  FDR DUMP SUCCESSFULLY COMPLETED FROM VOL=PROD01
FDR999  FDR SUCCESSFULLY COMPLETED
```

The backup created by this example is a normal FDR full-volume backup. There are no special considerations for restoring the entire disk volume or individual data sets from it. For example:

```
//RESTORE EXEC      PGM=FDR, REGION=2M
//SYSPRINT DD          SYSOUT=*
//SYSUDUMP DD          SYSOUT=*
//TAPE1     DD          DSN=BACKUP.VPROD01(0), DISP=OLD
//DISK1     DD          UNIT=SYSALLDA, VOL=SER=PROD01, DISP=OLD
//SYSIN     DD          *
      RESTORE TYPE=FDR
```

CONTINUED . . .

26.07 CONTINUED . . .

**FDRDSF DATA
SET BACKUP**

The Snapshot target volumes at addresses 01F2 and 01F3 are part of a pool of targets. In this run they will be used to snap volumes CICS01 and CICS02. When ready to execute, updates to the original volumes may need to be quiesced. The step SNAP will copy all data from the original volumes to the targets; when it completes, updates to the original volumes can begin again. Step BACKUP will backup the Snapshot copies of selected data sets on volumes CICS01 and CICS02, processing the two backups in parallel. After the backup is complete, FDR will release all of the back-end (internal disk) storage associated with the snapped volume.

```
// SNAP      EXEC      PGM=FDRSNAP
//SYSPRINT DD      SYSOUT=*
//TAPE1      DD      DUMMY
//SYSIN      DD      *
        SNAP TYPE=FDR
        MOUNT VOL=CICS01, SNAPUNIT=01F2
        MOUNT VOL=CICS02, SNAPUNIT=01F3
//BACKUP     EXEC      PGM=FDRDSF, REGION=2M, COND=(0, NE, SNAP)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1      DD      UNIT=SYSALLDA, VOL=SER=CICS01, DISP=OLD
//TAPE1      DD      DSN=BACKUP.VCICS01(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSPRIN1 DD      SYSOUT=*
//DISK2      DD      UNIT=SYSALLDA, VOL=SER=CICS02, DISP=OLD
//TAPE2      DD      DSN=BACKUP.VCICS02(+1), UNIT=TAPE, DISP=(, CATLG)
//SYSPRIN2 DD      SYSOUT=*
//SYSIN      DD      *
        DUMP TYPE=DSF, ATTACH, SNAP=(USE, REL)
        SELECT DSN=CICS**
```

The backups created by this example are normal FDRDSF data set backups. There are no special considerations for restoring data sets from them. For example:

```
//RESTORE EXEC      PGM=FDRDSF, REGION=2M
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//TAPE1      DD      DSN=BACKUP.VCICS01(0), DISP=OLD
//SYSIN      DD      *
        RESTORE TYPE=DSF, RECAT, VRECAT
        SELECT DSN=CICS3. LOADLIB
```

26.07 CONTINUED . . .

**FDRCOPY
FROM A
POINT-IN-TIME
IMAGE**

An SMS storage group consists of 2 volumes, SMS001 and SMS002, which have permanently assigned targets at addresses 01F4 and 01FA. You want to make a copy of data sets on those volumes but you can only quiesce updates for a few moments. The step SNAP will copy the online volumes to the offline target volumes. Step COPY will copy the selected data sets from the snapped volumes, renaming the data sets during the copy; the data sets will be copied to online volumes chosen by SMS as described in [Section 21](#). These copies might be used for parallel batch processing.

```
// SNAP      EXEC    PGM=FDRSNAP,REGION=0M
//SYSPRINT   DD      SYSOUT=*
//SYSIN      DD      *
  SNAP TYPE=FDR
  MOUNT VOL=SMS001,SNAPUNIT=01F4
  MOUNT VOL=SMS012,SNAPUNIT=01FA
//COPY      EXEC    PGM=FDRCOPY,REGION=0M,COND=(0,NE,SNAP)
//SYSPRINT   DD      SYSOUT=*
//SYSUDUMP   DD      SYSOUT=*
//DISK1      DD      UNIT=SYSALLDA,VOL=SER=SMS001,DISP=OLD
//DISK2      DD      UNIT=SYSALLDA,VOL=SER=SMS002,DISP=OLD
  COPY TYPE=DSF,SNAP=(USE,REL)
  SELECT DSN=PROD.PAYROLL.*,NEWI=..+COPY,STORCLAS=BATCH
```

**FDRCOPY
BETWEEN
SNAPSHOT
VOLUMES**

FDRCOPY can copy or move data sets to new volumes, and can rename and/or recatalog the output data sets. When FDR InstantBackup is installed, FDRCOPY is enhanced to automatically recognize that the input and output volumes for a given data set are in the same Snapshot-capable subsystem and automatically invokes the Snapshot hardware (if available) to do an instantaneous copy. When the volumes for a given data set are not in the same subsystem or Snapshot is not installed, normal read/write is used. This function is automatic, requiring no special operands, so the samples below are normal FDRCOPY jobs. If Snapshot was used for a given data set, its FDR311 message will contain the text "INSTANT".

This example copies data sets to new names, inserting a second-level index of COPY.

```
//COPYFILE   EXEC    PGM=FDRCOPY,REGION=2M
//SYSPRINT   DD      SYSOUT=*
//SYSUDUMP   DD      SYSOUT=*
//SYSIN      DD      *
  COPY      TYPE=DSF,MAXCARDS=1000
  SELECT    CATDSN=ABC.*,NEWI=..+COPY,
            NVOL=(TST001,TST002)
```

This example moves data sets to new volumes, recataloging them to their new location.

```
//MOVEFILE   EXEC    PGM=FDRCOPY,REGION=2M
//SYSPRINT   DD      SYSOUT=*
//SYSUDUMP   DD      SYSOUT=*
//SYSIN      DD      *
  MOVE      TYPE=DSF
  SELECT    CATDSN=MY.DATASET,NVOL=3390T1
  SELECT    CATDSN=*.DSNDBC.*.*.10001.A*,NVOL=3390T1
```

26.07 CONTINUED . . .

FASTCPK Normally FASTCPK must read every data track being moved, and write it to its new location. When SNAPSHOT=YES is specified and the volume being COMPAKed is in an disk subsystem with Snapshot support installed, the Snapshot facility is used to copy data tracks from one location to another internally. The data is not sent to or from the host. This usually results in a reduction in elapsed time. If the volume is not in a Snapshot-capable subsystem, or the volume contains DSORG=DA (Direct Access) data sets, normal read/write is done. This example will release space from all sequential, partitioned and VSAM data sets that have a secondary allocation quantity specified.

```
//FASTCPK EXEC PGM=FDRCPK,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSMAP DD SYSOUT=*
//SYSSUMM DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
        COMPAKT TYPE=FASTCPK,VOL=RVA001,SNAPSHOT=YES,
        CPKFREEEX=20,LOG=YES,NOSECOND=NORLSE,
        PSRLSE=ALL,PORLSE=ALL,VSRLSE=ALL
```

Note: if you simulate FAST COMPAKTION on a **Snapshot-capable** volume, the estimated elapsed time displayed by the simulation is calculated based on standard disk I/O times, not the reduced times provided by SNAPSHOT=YES. Your actual elapsed time with SNAPSHOT=YES will usually be much better than the simulation estimate.

26.08 EXAMPLES FOR ABR

**ABR FULL
VOLUME
BACKUP**

Snapshot target volumes have been permanently assigned to the online payroll volumes PAYxxx.

Step FULL1 will snap each online volume to its snapshot target to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
              SNAP   TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
              MOUNT  VOL=PAY001,SNAPUNIT=07C8
              MOUNT  VOL=PAY002,SNAPUNIT=07C9
              MOUNT  VOL=PAY003,SNAPUNIT=07CA
```

The SYSPRINT output of step FULL1 will look something like:

```
FDR304  ABR SNAP      REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07C8,DSNAME=FDRABR.VPAY001.C1000300,FILE=0000,VOL=SER=PAY001
FDR306  SPLIT        SUCCESSFULLY COMPLETED
FDR304  ABR SNAP      REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07C9,DSNAME=FDRABR.VPAY002.C1000300,FILE=0000,VOL=SER=PAY002
FDR306  SPLIT        SUCCESSFULLY COMPLETED
FDR304  ABR SNAP      REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07CA,DSNAME=FDRABR.VPAY003.C1000300,FILE=0000,VOL=SER=PAY003
FDR306  SPLIT        SUCCESSFULLY COMPLETED
```

The FDR304 message shows identifies this as a SNAP, and the FDR305 message identifies the snapshot target unit. It also shows the ABR backup data set name created, and that it was cataloged to the input volume.

The SYSPRINx output will also contain a message to indicate that a SNAP was done:

```
FDR218   FDRSNAP UNIT=07C0 IS SNAPPED FROM UNIT=07C8
```

CONTINUED . . .

26.08 CONTINUED . . .

**ABR FULL
VOLUME
BACKUP
(continued)**

Step FULL2 will do the actual backups. SNAP=(USE,REL) tells ABR to determine if a snapped point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR remembers the unit addresses of the snapped offline volumes specified in the FULL1 step; there is no need to specify them here. DCT=YES requests that the HSDM (High Speed Data Mover) hardware option be used if it is installed on your disk subsystem; HSDM can reduce backup times up to 60%.

```
//FULL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11     DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2      DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22     DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN      DD     *
              DUMP TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,SNAP=(USE,REL),DCT=YES
              MOUNT VOL=PAY*
```

The SYSPRINT output for step FULL2 will be similar to:

```
FDR219 FDR VOL=PAY001 IS BEING DUMPED FROM OFFLINE UNIT=07C8
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY001.C1000300,FILE=0005,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY001.C2000300,FILE=0001,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY002 IS BEING DUMPED FROM OFFLINE UNIT=07C9
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY002.C1000300,FILE=0012,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY002.C2000300,FILE=0001,VOL=SER=B15600
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY003 IS BEING DUMPED FROM OFFLINE UNIT=07CA
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY003.C1000300,FILE=0006,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY003.C2000300,FILE=0002,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR224 FDR VOL=PAY004 IS BEING DUMPED FROM ONLINE UNIT BECAUSE SNAP/SPLIT WAS NOT ATTEMPTED
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY004,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY004.C1000300,FILE=0013,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY004.C2000300,FILE=0002,VOL=SER=B15600
FDR306 DUMP SUCCESSFULLY COMPLETED
```

FDR219 messages indicate that a previously created point-in-time image was found for the volume being dumped. FDR305 messages show the actual location of the backup data sets. FDR218 messages indicate that the REL option caused ABR to release the back-end storage associated with the snapshot target volume, to reduce the NCL (Net Capacity Load) of the subsystem.

In this example, volume PAY004 was not snapped. The FDR224 message indicates that there is no point-in-time image of the volume so that a normal ABR backup will be done. If you intended to do point-in-time backups all of the volumes in this backup, you must do a SNAP for PAY004 in the FULL1 step; otherwise the backup from PAY004 will not be taken at the same point-in-time as the other volumes. If you have no need to do point-in-time backup on PAY004, you can ignore the message. So, you can include in this ABR step volumes which have no point-in-time requirements; they will be backed up normally.

Note that if you attempted to SNAP a volume but the SNAP failed for some reason (perhaps because the volume is not in a Snapshot-capable subsystem), ABR will set a flag in the ABR Model DSCB on the online volume. The DUMP step will fail the backup of that volume with message FDR211 and a non-zero return code, assuming that you intended it to be a point-in-time backup.

26.08 CONTINUED . . .

**ABR
INCREMENTAL
BACKUP**

The snapshot target volumes at addresses 01F2 and 01F3 are part of a pool of targets. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step INCR1 will snap each online volume to create a point-in-time image and mark the backup as complete. It will create a new ABR backup cycle and update the online volume with information about the new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that the RETPD= parameter must be specified the same as in the INCR1 step which follows.

```
// INCR1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY
//SYSIN       DD     *
              SNAP  TYPE=ABR,DSNENQ=USE,RETPD=14,ENQERR=NO
              MOUNT VOL=CICS01,SNAPUNIT=01F2
              MOUNT VOL=CICS02,SNAPUNIT=01F3
```

Step INCR2 will do the actual backups. SNAP=(USE,REL) tells ABR to determine if a snapped point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that remembers the addresses of the offline snap targets; there is no need to specify them here.

```
// INCR2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP)
//SYSIN       DD     *
              DUMP  TYPE=ABR,DSNENQ=USE,DATA=USED,RETPD=14,ENQERR=NO,SNAP=(USE,REL)
              MOUNT VOL=CICS01
              MOUNT VOL=CICS02
```

For examples of messages and other notes, see the preceding example, "ABR FULL VOLUME BACKUP"

26.08 CONTINUED . . .

ABR RESTORE If it is necessary to do a full-volume restore of one of the volumes involved in the point-in-time backup before the backup is moved to tape, ABR will do it automatically. No special parameters or JCL is required. ABR will restore from the point-in-time image or from tape automatically.

For example, this will restore the four volumes in the previous full-volume backup example:

```
//REST      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT  DD     SYSOUT=*
//SYSIN     DD     *
RESTORE    TYPE=FDR,DYNTAPE,CONFMESS=NO,ONLINE,CPYVOLID=YES
SELECT     VOL=PAY001
SELECT     VOL=PAY002
SELECT     VOL=PAY003
SELECT     VOL=PAY004
```

Volumes PAY001-003 will be restored from the point-in-time images on the snap target volumes unless they have been moved to tape. Volume PAY004, since it has no point-in-time image, will always be restored from tape.

However, there are some considerations:

- the point-in-time volume image is always a full-volume image, even if a incremental backup was done. If ABR discovers that the first ABR backup cycle to be restored is a point-in-time image, it simply uses the appropriate hardware features to quickly copy the entire volume back to the target online volume and does not need to read any further backups from tape. Since full-volume restores typically start with the most recent backup of the volume, this provides the fastest recovery.
- but if the most recent ABR backup cycle is on tape, but another lower-numbered cycle is a point-in-time image, ABR will bypass restoring from that image (with a message); any data sets whose most recent backup was on that image may be corrupted or back-level. This could occur if you did a SNAP to create a cycle but never did a DUMP to move it to tape, and followed that with a regular ABR backup direct to tape. This is not normal and should not occur if your FDR InstantBackup procedures are setup correctly.
- If you are restoring to a volume other than the original volume (NVOL= on the SELECT statement) and the target volume is not in the same disk subsystem as the point-in-time image, ABR will do a normal disk-to-disk copy to copy all tracks in the image to that volume.
- If ABR attempts to restore from a point-in-time image (ABR backup data set cataloged to the original volume), but that image no longer exists, ABR will bypass that backup. If it was an incremental backup, it will restore the rest of the cycles in that generation. If it was a full-volume backup (cycle 00), the restore will fail (you can uncatalog the backup or specify GEN= and CYCLE= to restore from an earlier generation). This can occur if multiple SNAP operations were done without an intervening DUMP to move the backup to tape. It can also occur at a disaster site if the restored ABR catalog is not consistent with the tapes that were brought off-site (some backup entries still indicate the point-in-time backup).

THIS PAGE WAS INTENTIONALLY LEFT BLANK

27.01 FDR INSTANTBACKUP FOR HITACHI DATA SYSTEMS SHADOWIMAGE

**FDR
INSTANTBACKUP
AND HDS
SHADOWIMAGE™**

FDR InstantBackup is an additional cost FDR facility for non-disruptive backup. It works in concert with ShadowImage, a licensed software feature of the Hitachi Data Systems (HDS) 7700E disk subsystem, and also with HS-DataPlex, a feature of the HDS 7700 disk subsystem, which allows for the "instant" creation of exact duplicates of existing DASD volumes.

FDR InstantBackup enhances:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY data set copy
- FDRABR volume backup (full-volume and incremental backup)

It provides non-disruptive backup of offline MVS DASD volumes, both SMS and non-SMS, using technology developed by Innovation Data Processing.

Note: the FDRABR support for ShadowImage uses an IBM API (Application Program Interface) for PPRC which was introduced in early 1998. The simplest way to tell if this API is installed is to check SYS1.MACLIB for a macro called ANTRQST. To get this API, if you are running

OS/390 V1.1-V1.3 (DFSMS 1.3) you must install IBM PTF UW46075

OS/390 V2.4-V2.5 (DFSMS 1.4) you may need to install IBM PTF UW46076

OS/390 V2.6 (DFSMS 1.4) the API is installed

OS/390 V2.7 and above (DFSMS 1.5 and above) the API is installed.

You may also need to install the appropriate PTF for IBM APAR OW43598 to OS/390 V1.1-V2.9 systems to correct a problem where IBM was incorrectly checking the results of a PSUSPEND command in the API.

SHADOWIMAGE

Traditional backups may disrupt normal operations. It obviously takes time to backup a data set or volume. If data sets are being updated by some application while they are being backed up, the backup may not be valid and the data sets may not be usable when they are restored. The usual response to this problem is to quiesce all updates to the data during the backup, disrupting normal operation until the backup is complete.

When used with Hitachi Data Systems' ShadowImage feature, FDR InstantBackup allows you to instantly freeze a copy of an online disk volume and easily backup or copy data from that frozen copy as if it was coming from the original volume. Updates to the data need to be quiesced only for the few moments necessary to create the frozen copy.

Note: although this documentation talks only about ShadowImage, FDR InstantBackup also supports Hitachi Data Systems Synchronous Remote Copy which can be used on 7700 classic systems as well as 7700E systems. These features, under different names, may also be available on disk systems supplied by Compaq in Europe.

**FDR
INSTANTBACKUP
ENHANCES FDR
COMPONENTS**

FDR InstantBackup enhances FDR, FDRDSF, and FDRCOPY to read from an offline copy of an online volume which was created as a point-in-time frozen image using ShadowImage. This creates a backup or copy of the volumes or data sets that looks exactly like a normal backup or copy, except that the data is frozen in time.

If you are licensed for FDRABR, FDR InstantBackup enhances ABR Volume Backups, both full-volume (TYPE=FDR) and incremental (TYPE=ABR/DSF/AUTO) to use ShadowImage, by creating "instant" frozen copies of the volumes to be backed up, and then moving that captured data to tape.

27.02 SHADOWIMAGE OPERATION

**HITACHI
DATA
SYSTEMS
VOLUME
IMAGING
OPTIONS**

The Hitachi Data Systems 7700/7700E storage subsystems are RAID implementations, providing automatic protection against data loss without special features. However, they also include two features which enhance this protection and offer other benefits such as the creation of test copies, and use of FDR InstantBackup.

- "Internal Remote Copy" allows creation of a second copy of a volume in the same subsystem with its own MVS address using Hitachi Synchronous Remote Copy (PPRC compatible). It requires an external ESCON cable. FDR (but not ABR) supports Internal Remote Copy. This is supported on both the 7700 and 7700E subsystems.
- The ShadowImage feature, allows up to 3 PPRC copies of an original volume in the same subsystem. No external cable is required. This document will refer only to ShadowImage but FDR will also work with the Internal Remote Copy. This licensed software feature is available for the 7700E.

Both features will automatically mirror any updates to the online volume to all of its copies, keeping all of them in synchronization with the online volume. The ShadowImage volumes:

- have their own device addresses
- are not permanently assigned to any online volume
- can be established to a online volume as a copy of any online volume in the disk subsystem with the same device format (3390 or 3380) and size. When initially established, all data on the online volume will be internally copied to the secondary volume.
- can be detached from its current online volume by totally breaking the connection (resetting to simplex mode). At this point the secondary volume is a frozen point-in-time image of the online volume and can be read or written, but no connection to the online volume is maintained; if it is again established to the online volume, the entire contents of the online volume must be recopied to the secondary.
- can be split from its current online volume by suspending the copy. At this point the secondary volume is a frozen point-in-time image of the online volume and it can be read (Internal Remote Copy) or read and written (ShadowImage). While suspended, the control unit records which tracks on the online volume and the split secondary volume have been updated; when the pair is resynchronized, only those updated tracks are copied to the secondary, bringing the volumes in synchronization much more quickly.

Under MVS, you also use the standard IBM PPRC (Peer-to-Peer Remote Copy) commands to manage ShadowImage volumes, including:

- the TSO commands CQUERY, CESTPAIR, CSUSPEND and CDELPAR
- the ICKDSF PPRCOPY subcommands QUERY, ESTPAIR, SUSPEND and DELPAIR

ShadowImage volumes do not replace normal RAID-1 or RAID-5 volumes but provide a RAID protected copy of online volumes. A given online volume can be duplexed to a ShadowImage RAID protected copy in either a RAID-1 or RAID-5 configuration independent of the RAID configuration the online volume is running under.

Your Hitachi Data Systems Account team can provide you with more information on ShadowImage

27.02 CONTINUED

**SHADOWIMAGE
WITHOUT FDR
INSTANT BACKUP**

Although ShadowImage can quickly create a point-in-time image of a production disk volume, traditional backup products can not directly use the ShadowImage volume image without additional considerations.

For example, a conventional backup product will require that the detached ShadowImage volume be brought online. Since MVS will not let you put two volumes with the same volume serial online, the ShadowImage volume serial must be changed.

Standard operating system utilities (e.g., ICKDSF) provide a facility for re-labeling a ShadowImage volume after it is detached, but re-labeling a volume invokes additional requirements for SMS managed volumes, VSAM clusters and cataloged data sets. The data sets on the re-labeled volume will appear to be uncataloged, which may cause problems during backup and restore unless they are renamed and recataloged.

A backup of a re-labeled volume will appear to be a backup from the new volume serial, not the original volume serial, so additional procedures are required to document this correspondence (e.g., backup of volume B is really a backup of volume A). Restore procedures must be modified to restore the backups to the correct volume serial (e.g., restore volume B back to volume A and re-label it as A).

Hitachi Data Systems also supports a software utility to rename and recatalog the data sets on the ShadowImage volume after the split. Although this could be used to backup those data sets with a standard backup utility, the installation must now use special procedures to restore each data set from its temporary name on the backup to its true name on the online volume. For large numbers of data sets, this could be onerous.

27.03 THE FDR INSTANTBACKUP SOLUTION**THE FDR
INSTANTBACKUP
SOLUTION**

FDR InstantBackup allows you to create and use a duplicate ShadowImage volume image without changing its volume serial or bringing it online. FDR full volume backup, FDRDSF data set backup and FDRCOPY data set copy use FDR InstantBackup technology to read the detached ShadowImage volume even while it remains offline. A non-disruptive backup or copy can start as soon as the ShadowImage volume is detached. Updates to the online volumes continue while you create the point-in-time backup or copy of the data.

For FDRABR, FDR InstantBackup makes volume backups a 2-step process. The first step executes a PSPLIT statement under FDRABR, to split previously established ShadowImage copies of the selected volumes. If the copy is fully synchronized with its online volume, this takes only seconds per volume. The split copy is actually a valid backup at that point-in-time; you can do a restore from it if recovery is required. The second step executes a normal FDRABR full-volume or incremental backup of the same volumes, except that ABR will read the offline point-in-time image and move it to a backup on tape. There is no need to identify the ShadowImage device to ABR; it knows the ShadowImage volume address of any given online volume.

You invoke FDR InstantBackup for FDR, FDRDSF, and FDRCOPY with standard FDR JCL which points to the online volume but which contains a special data set name that identifies the offline ShadowImage volume just split from the online volume. FDR InstantBackup verifies that the ShadowImage device is offline and that its volume serial matches the serial of the online disk. It then will read the offline disk in exactly the same way that it would read the online disk. FDR InstantBackup produces a backup that appears to be a conventional backup of the online volume serial. FDR InstantBackup blends its unique off-line, high-speed non-disruptive backup together with a traditional restore complete with FDR's powerful logical file capability.

There are no special considerations for restores from a backup created by FDR InstantBackup. You can restore entire volumes or individual data sets from those backups. The target volumes for restore will be online volumes, not the ShadowImage volumes. Since the backups created by FDR InstantBackup appear to be backups of the online volume, there are no special volume serial or data set name concerns.

When FDR InstantBackup is used with FDRCOPY, you can freeze a ShadowImage volume image of the online volumes containing the desired data sets, and then create copies of those data sets on other online volumes. Data sets can be renamed and cataloged during the copy.

27.04 SETTING UP ShadowImage VOLUMES

**SHADOWIMAGE
VOLUME
DEFINITION**

To use ShadowImage volumes you must have the proper features installed on your 7700/7700E Subsystem, and you must have sufficient disk capacity within the subsystem to support the newly created ShadowImage volumes. You may also need to ask your Hitachi Data Systems CS&S Support Team to enable "special option 20" which allows read access to suspended Hitachi Remote Copy volumes (not required for ShadowImage).

To determine the number of ShadowImage volumes to be created, you must identify all of the online volumes for which you intend to use FDR InstantBackup and categorize them by type (3380 or 3390) and size (number of cylinders). A ShadowImage volume can be only be paired with a online volume of the same type and equal to or larger size; for use with FDR InstantBackup, the size (number of cylinders) should be equal. ShadowImage volumes can be any offline volume within the same 7700/7700E Subsystem as the online volume.

If you have sufficient capacity, you can assign ShadowImage volumes for every online volume involved. If not, you can use a "dynamic pool" of ShadowImage volumes which can be used to create duplicates of online volumes as needed.

The device addresses assigned to the ShadowImage volumes, in a 7700/7700E Subsystem, must be defined as normal 3380 or 3390 disks in your I/O configuration (IOCP or HCD). However, as shown later in this section, when a ShadowImage volume is detached from its online volume for use with FDR InstantBackup it will have the same internal volume serial as the online volume. This is normally not a problem since it will be offline to MVS, but it can cause confusion if you IPL your MVS system since the IPL process will detect that two volumes which have the same serial and ask the operator which should be left offline (message IEA213A). Obviously, the real online volume should remain online so the operator must respond with the ShadowImage volume address. **To avoid the possibility of the operator replying incorrectly, you may want to mark the ShadowImage volumes as OFFLINE in your I/O configuration, so that MVS will not attempt to access them at IPL time.** FDR InstantBackup will be able access them even though they are offline. However, if you use ShadowImage only with FDR InstantBackup for ABR, this is not necessary since ABR will modify the volume label to prevent the volume copy from being varied online.

View Device Parameter / Feature Definition				Row 1 of 6
Command ==>	-----	Scroll ==>	PAGE	
Configuration ID . . .	CURPI099	Common Configuration data		
Device number . . .	01E0	Device type . . .	3390	
Generic / VM device type . . .	3390			
Parameter/	Value	Req.	Description	
Feature				
OFFLINE	Yes		Device considered online or offline at IPL	
DYNAMIC	Yes		Device supports dynamic configuration	
LOCANY	Yes		UCB can reside in 31 bit storage	
ALTCTRL	No		Separate physical control unit path	
SHARED	Yes		Device shared with other systems	
SHAREDUP	No		Shared when system physically partitioned	

For example, in HCD, the MVS options for a ShadowImage device will look similar to:

In IOCP/MVSCP, specify OFFLINE=YES on the IODEVICE macros for the ShadowImage devices. The following is an example of a 7700/7700E attached via 2 ESCON channels (CHPIDs 40 and 41), with 128 online disk addresses (300-37F) and 64 ShadowImage addresses (380-3BF):

```
CNTLUNIT CUNUMBR=008,PATH=(40),LINK=(C0),
        UNIT=3990,UNITADD=((00,192))
CNTLUNIT CUNUMBR=009,PATH=(41),LINK=(C4),
        UNIT=3990,UNITADD=((00,192))
IODEVICE ADDRESS=(300,128),CUNUMBR=(008,009), ** ONLINE
        UNIT=3390,UNITADD=00,FEATURE=SHARED
IODEVICE ADDRESS=(380,64),CUNUMBR=(008,009), ** ShadowImage
        UNIT=3390,UNITADD=80,FEATURE=SHARED,OFFLINE=YES
```

Note that some IOCP documentation may indicate that UNITADD=00 is required for ESCON; what it really means is that the value for the **first** IODEVICE definition for a given control unit must be 00. So when varying characteristics must be specified for some devices on a control unit (such as OFFLINE=YES), you can have multiple IODEVICE statements with varying values for UNITADD.

27.04 CONTINUED

**SHADOWIMAGE
VOLUMES WITH
FDR
INSTANTBACKUP**

A ShadowImage volume is paired with an online volume using the TSO command
 CESTPAIR DEVN(X'0C84') PRI(X'0080',30158,X'04') SEC(X'0080',30158,X'1A')
 or the ICKDSF command:

```
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'04')
SEC(X'0080',30158,X'1A')
```

The "triplet" of identifying information in PRI (the online volume) and SEC (the ShadowImage copy) can be obtained for any given HITACHI DATA SYSTEMS disk address by the TSO command CQUERY DEVN(X'uuuu').

The ESTPAIR will synchronize the ShadowImage volume with the online volume by copying all data on the online volume to the ShadowImage volume; this will take several minutes per volume to be copied, depending on physical size and subsystem activity. Once the copy is complete, all updates to the online volume will be reflected on the ShadowImage volume, **the 7700/7700E Subsystem hardware insures that the ShadowImage volume is always an exact copy of its online volume.** While the ShadowImage volume pair is active, in duplex mode, the host cannot directly address the ShadowImage volume.

Since **all** data on the online volume is mirrored on the ShadowImage volume, including the volume label and the VTOC, VTOCIX, and VVDS, the ShadowImage volume has the same volume serial as the online volume.

At the point in time where a backup or copy of the data on the online volume is desired, the ShadowImage volume can be either split from the online volume by a PPRC SUSPEND command or the pair can be reset to simplex mode by a PPRC DELPAIR command. The ShadowImage volume will be an exact copy of the online volume at the time that the suspension or reset occurred. The split of the ShadowImage volume permits it to be re-synchronized with the online volume once the backup is completed. The reset to simplex removes any logical hardware association of the ShadowImage volume to the online volume and permits the ShadowImage volume to be used in a different ShadowImage pair once the backup is complete.

For FDRABR, you split the ShadowImage volume using the PSPLIT statement of FDRABR. This allows ABR to identify the address of the ShadowImage volume for later processing. See the examples later in this section for more details.

For FDR, FDRDSF, and FDRCOPY, you split the ShadowImage volume using either the TSO command:

```
CSUSPEND DEVN(X'0C84') PRI(X'0080',30158,X'04') SEC(X'0080',30158,X'1A')
```

or the ICKDSF command:

```
PPRCOPY SUSPEND DDNAME(DISK1) PRI(X'0080',30158,X'04') SEC(X'0080',30158,X'1A')
```

Once split or reset is complete, the ShadowImage volume can be accessed by its assigned device address and can be backed up or copied with FDR InstantBackup. FDR InstantBackup can backup or copy data from the ShadowImage volume even though it is offline and has a serial that duplicates the online volume.

When the backup or copy is complete, you can re-synchronize a split (suspended) ShadowImage volume with the online volume PPRC ESTPAIR MODE(RESYNC) command. During the time when the ShadowImage volume pair was split, the 7700 Subsystem keeps track of all data updated on the online volume. Re-synchronization copies only the updated data from the online volume to the ShadowImage volume, so it is much faster than a full establish operation. Once the ShadowImage volume is again synchronized with the online volume, it is ready to be used for another FDR InstantBackup operation.

For FDRABR, the ABR step which moves the backup data to tape has an option PPRC=(USE,RET) which will automatically re-synchronize the ShadowImage volume when the backup is complete.

27.04 CONTINUED

For FDR, FDRDSF, and FDRCOPY, you re-synchronize the ShadowImage volume using either the TSO command:

```
CESTPAIR DEVN(X'0C84') PRI(X'0080',30158,X'04')
SEC(X'0080',30158,X'1A') MODE(RESYNC)
```

or the ICKDSF command:

```
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'04') -
SEC(X'0080',30158,X'1A') MODE(RESYNC)
```

Note that the copying of data to ShadowImage volumes is done internally by the disk subsystem and is a low priority task (real I/Os get priority). However, the amount of time required is not really important unless you need to detach the ShadowImage volume and do a backup before the volumes are completely synchronized.

You never need to bring a ShadowImage volume online or re-label it or re-initialize it with ICKDSF, as long as it is used only for FDR InstantBackup.

**SETTING UP A
SHADOWIMAGE
ENVIRONMENT**

There are 2 ways to setup ShadowImage use:

- 1) you can permanently assign ShadowImage volumes to every online volume for which you plan to use FDR InstantBackup backup or copy. This will require more 7700/7700E capacity and more device addresses, but is much simpler to administer. The ShadowImage volume is actively mirroring it's online volume all the time, except during the backup or copy, so the duplex pair is established only once. The pair is split before the backup or copy. Afterwards, the pair will quickly re-synchronize the volumes since only changed data is copied.
- 2) you can create a pool of ShadowImage volumes and assign them to online volumes as necessary. This requires less 7700 capacity and fewer device addresses, but is more complex to administer. You will have to activate the duplex pair every time it is to be used. It will take additional time to copy all of the data from the online volume, especially if multiple pairs are being activated. You will require operational procedures to insure that two FDR InstantBackup jobs don't try to use the same ShadowImage volume for two different online volumes at the same time. You will reset the pair to simplex mode before the backup/copy.

You can use a combination of the two, some online volumes with permanent ShadowImage volumes (perhaps for volumes with frequent or time-critical backups) and some with ShadowImage volumes assigned from a pool (perhaps for less frequent or non-time-critical backups). The pool of volumes set aside for use with ShadowImage are available for normal use during the time the ShadowImage capability is not required. The pool consists of any offline volume installed in the 7700/7700E subsystem.

In either case, the ShadowImage volumes must match the size (number of cylinders) and format (3380 or 3390) of the online volumes to which they will be assigned.

CONTINUED . . .

27.05 FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY

FDR InstantBackup uses standard FDR JCL with one modification. (see "FDR INSTANT BACKUP FOR ABR" in [Section 27.06](#) for FDRABR requirements). The FDR control statements used are standard except for a few considerations. For complete details on other JCL and control statement requirements of FDR, FDRDSF, and FDRCOPY, please see [Sections 10, 20, and 21](#) of the FDR manual, respectively.

FDR InstantBackup uses a special convention for the JCL statements that identify the disk to be backed up or copied, i.e. DISKx DDs. As shown in the examples below, the DD points to the online volume but includes a special data set name which tells FDR the device address of the ShadowImage volume to be read in place of the online volume. FDR will verify that the volume serial of the online volume is the same as that of the offline ShadowImage volume.

These special DISKx DDs can be used with:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY input disks for COPY operations only. Although FDRCOPY can copy data sets without the use of DISKx DD statements, DISKx DDs are required to cause FDRCOPY to read the offline ShadowImage volume. You can still select data sets normally, including selecting them from the catalog, but FDR InstantBackup will be used for any input volume identified on a DISKx DD statement.

They can also be used for FDR full-volume copies (COPY TYPE=FDR) but in most cases it would be more useful to copy the online volume directly.

FDRCOPY MOVE operations from a point-in-time copy are not supported by FDR InstantBackup, since it deletes the input data set. You should always move the live, online version of a data set.

**NON-SMS
MANAGED
VOLUMES**

For non-SMS volumes, a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,DISBAND,UNIT=3390,
//          VOL=SER=vvvvvv,DISP=OLD
```

The UNIT and VOLSER (vvvvvv) must identify the online disk volume. Note that the FDR.USE.UNITuuuu.DISBAND data set does not really exist anywhere in your system, but the presence of the special data set name on the DD statement is very important. It is this name that directs FDR to access the offline ShadowImage volume device located at UCB address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address. A leading zero must be added if it is a 3-digit address, e.g., DSN=FDR.USE.UNIT01F3.DISBAND. This special data set name is never opened so it will cause no problems with security. The ".DISBAND" at the end of the name is optional, but is recommended for Hitachi disks; it disbands the path group established by FDR InstantBackup so that PPRC can again use the target volume.

27.05 CONTINUED

**SMS-
MANAGED
VOLUMES**

For SMS-managed volumes a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,DISBAND,DISP=OLD
```

Because MVS will ignore a volume serial in the JCL if it points to a SMS-managed volume, the special FDR.USE.UNITuuuu.DISBAND name must be cataloged. **This data set does not really have to exist on the original volume, but it must be cataloged to that volume! Just as for a non-SMS volume, its presence on the DD statement is very important.** It is this name that directs FDR to access the ShadowImage volume on the offline device located at device address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address; add a leading zero if necessary.

Here is one technique for creating the catalog entry required for FDR InstantBackup. This step will catalog the special FDR.USE.UNITuuuu.DISBAND name to an SMS managed volume. It could be run immediately before a backup.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE FDR.USE.UNITuuuu.DISBAND NOSCRATCH
DEFINE NONVSAM(NAME(FDR.USE.UNITuuuu.DISBAND) -
DEVICETYPE(33xx) VOLUME(vvvvvv))
```

The DEVICETYPE and VOLUME parameters identify the SMS-managed online volume. The special FDR.USE.UNITuuuu.DISBAND name will direct FDR to access the offline ShadowImage volume device located at device address uuuu instead of the original online SMS managed online volume.

**FDR CONTROL
STATEMENTS**

The control statements used with FDR InstantBackup are the normal statements documented for FDR ([Section 10](#)), FDRDSF ([Section 20](#)), and FDRCOPY ([Section 21](#)) with these modifications:

- Since FDR InstantBackup is not reading the live data, it is never appropriate to use the ENQ= or DSNENQ= operands to enqueue on the VTOC of the ShadowImage volume or the data sets on the ShadowImage volume. To do so would unnecessarily prevent access to the live data on the online volume.
- For FDRCOPY, specify DSNENQ=NONE. Also, the input volumes must be specified by DISKx DD statements (to identify the ShadowImage volume addresses) so you should not select data sets with CATDSN= or VOL=.
- Do not use the FDRCOPY MOVE statement when the input is a ShadowImage volume. If you do so, the live data sets on the online volume will be deleted at the end of the move.

CONTINUED . . .

27.05 CONTINUED

**FDR
INSTANTBACKUP
OPERATION**

If you have a permanently assigned ShadowImage volume for every online volume, use these steps:

- 1) Execute (C)ESTPAIR command, one time, to create the duplex pair and assign the ShadowImage volume to its online volume. Be sure to do this far enough ahead of your first backup to allow the volumes to be synchronized.
- 2) When ready to backup or copy data, quiesce system and/or application update activity on the online volumes, if required, and execute (C)SUSPEND command to split each pair required. Once the split ends, you can re-enable updates on the online volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with the special DISKx DD statements described above to access the ShadowImage volume.
- 4) When you no longer need the "frozen" data, execute the (C)ESTPAIR MODE(RESYNC) command to resynchronize the ShadowImage volume with its online volume. Unless there has been a lot of update activity on the online volume, this will complete quickly.

If you are using a ShadowImage volume pool (described earlier), use these steps:

- 1) Execute the (C)ESTPAIR command for each online volume to be backed up, attaching it to one ShadowImage volume from the pool. Be sure to define the pair far enough in advance of the required backup time to allow the synchronization to complete.
- 2) When ready to backup or copy data, quiesce update activity on the online volumes, if required, and execute (C)DELPAR command to detach the ShadowImage volume (reset to simplex). Once the detach ends, you can re-enable updates on the online volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with the special DISKx DD statements described above to access the ShadowImage volumes.
- 4) There is no need to re-establish the duplex pair unless you plan to do another FDR InstantBackup of the same online volume. Once the backup is complete the ShadowImage volume is available for use with a different online volume. However, you will need operational procedures to insure that different FDR InstantBackup jobs do not attempt to use the same ShadowImage volume at the same time.

For details on the proper commands to control ShadowImage operations, please see Hitachi Data Systems documentation or consult with your Hitachi Data System account team.

**MULTIPLE
SYSTEMS**

FDR InstantBackup requires that both the original volume and its offline ShadowImage copy be accessible from the MVS system where the backup is run. If you execute the split on one system, but execute the backup on a second MVS system which does NOT have access to the online volume, FDR InstantBackup will fail. There is a circumvention for FDR and FDRDSF, but you must contact Innovation for assistance in implementing it.

27.06 FDR INSTANTBACKUP FOR ABR

When a normal ABR full-volume or incremental backup completes successfully, it is recorded by doing 2 things:

- ABR catalogs the backup data set it just created in the ABR catalog
- ABR updates the VTOC on the volume to record in the ABR Model DSCB the generation and cycle of the backup just created, and updates the current backup information for every selected data set

So, as soon as the backup data set is complete and available, all ABR records point to it and restores can be immediately requested.

**OPERATION
OF FDR
INSTANTBACK
UP FOR ABR**

With FDR InstantBackup for ABR, all of the same information must be recorded on the online volume, to document the backup. However, these changes can't wait until the backup of the offline point-in-time image to tape is complete, since changes may have occurred to the online volume in the meantime (such as updating a data set again, deleting or allocating data sets).

To solve this, when ABR is invoked to capture the point-in-time image, it will update the online volume immediately. The ABR generation and cycle of the next backup is determined and all of the steps shown above are taken, except that the backup is cataloged to the online volume itself (something that would never occur in a normal ABR backup) so this tells ABR that the backup is a point-in-time image. This functions as a true backup, since you can do restores from it.

Later, ABR will process the point-in-time image and create a normal ABR backup from the offline image, exactly like it would do the backup of the online volume if InstantBackup were not involved. It will recatalog the backup data set to the tape just created. At this point, the backup looks exactly like a normal ABR backup of the online volume.

Even before the backup is moved to tape, ABR can restore the entire volume from the point-in-time image.

**EXECUTING FDR
INSTANTBACKUP
FOR ABR**

To use FDR InstantBackup for ABR, you must divide your existing ABR volume backup jobs into two steps which may be executed in one job or in two separate jobs:

- 1) the first step will capture all the point-in-time images of the selected online volumes using the ABR statement PSPLIT. This will normally take only a few seconds per volume selected. We recommend that updates to the online volumes be suspended as far as possible during this step. As soon as it finishes, updates to the online volumes can begin again, since the point-in-time captured image of the volume exists. To insure that the snapped volume is not accidentally brought online, ABR will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail.
- 2) the second step will create the actual backup tapes, reading the captured point-in-time images. It will be identical to the original ABR step, except that ABR operand PPRC= must be added to invoke this support. This ABR step can also include backups of volumes which do not have captured images; these volumes will be backed up normally (with a warning message). It also optionally releases the captured images when it is done with them by re-synchronizing the ShadowImage volume to its online volume.

When you split the ABR step into two, each step should contain the same JCL and ABR control statements as the original, with a few changes. It is vital that all ABR operands be identical in both steps, especially the TYPE= operand, except as noted below (actually, the second step will fail if the TYPE= operand is different). The examples later in this section illustrate the changes required.

27.06 CONTINUED

EXECUTING FDR
INSTANTBACKUP
FOR ABR
(CONTINUED)

In the first step:

- replace the operation name DUMP on the first ABR control statement with PSPLIT
- change all TAPE_x and TAPE_{xx} DD statements to DUMMY. **However, if the original TAPE_x/xx DD statements contained RETPD= or EXPDT= operands, include the same operands on the DUMMY DD statements so that ABR can record the actual expiration of the backups. Similarly, if you coded the ABR operand RETPD= on the DUMP statement, include it on the PSPLIT statement as well.**
- Innovation recommends the use of DSNENQ=USE and ENQ=RESERVE so that ABR can properly protect the volume and data sets during the few moments necessary to create the point-in-time image. However, unlike normal ABR operations, where ABR will still backup a data set or volume even if DSNENQ= finds active data sets, PSPLIT will not create the point-in-time if it finds data sets in use, which will also cause the second step to fail trying to backup that image. You can override this by specifying the ENQERR=NO operand on the DUMP statement, or by adding statements such as:
 SELECT DSN=dsname,DSNENQ=NONE
 for data sets which you expect to be active at the time of the PSPLIT; this is supported even for PSPLIT TYPE=FDR
- If the backup includes other volumes in other disk subsystems, or volumes for which you simply don't want to do a point-in-time backup, they should be omitted from this step, even though they can be included in the second step
- You may need to make one more change to the first step. In ABR, you identify the volumes to be processed with DISK_{xxxx} DD statements or ABR MOUNT statements; MOUNT statements can specify groups of volumes, e.g., VOL=TSO*. For each online volume to be processed for PSPLIT, ABR can usually identify the MVS address of the secondary ShadowImage volume assigned to it. However, there are two circumstances where you may need to tell ABR the ShadowImage address:
 - 1) the online and ShadowImage volumes are not in the same range of MVS addresses. For example, the online volume is in a logical control unit with addresses X'17C0-17FF' but the ShadowImage volume is in a LCU with addresses X'1840-187F'. ABR may still be able to identify the ShadowImage volume but it will take more time, so specifying the ShadowImage address will save time.
 - 2) If you have more than one ShadowImage volume or PPRC copy paired to the online volume, you **must** identify the ShadowImage address to be used.

To identify the ShadowImage addresses, you will have to change the ABR input to something similar to:

```
PSPLIT TYPE=ABR
MOUNT VOL=vvvvv1,PPRCUNIT=uuu1
MOUNT VOL=vvvvv2,PPRCUNIT=uuu2
```

with one MOUNT statement for every ShadowImage volume. You do not have to make the same change in the second step; it can continue to select volumes with MOUNT or DISK_{xxxx} DD statements.

27.06 CONTINUED

**EXECUTING FDR
INSTANTBACKUP
FOR ABR
(CONTINUED)**

In the second step:

- add to the DUMP statement the operands

PPRC=USE (read point-in-time image) or

PPRC=(USE,RET) (read point-in-time image and resync copy when done)

PPRC=USE is used when you intend to pair the ShadowImage volume with a different online volume after the backup is complete; PPRC=(USE,RET) is used when the ShadowImage volume is permanently assigned to the online volume.

- DSNENQ= and ENQ= will be ignored by ABR when dumping point-in-time volume images, since these would prevent access to the online data sets and prevent changes to the online volume's VTOC, not the offline point-in-time image. However, they will be honored if you backup other non-point-in-time volumes in the same step

Since the target volumes can hold only one point-in-time image at a time, you must be sure that another ESTPAIR is not done on the split ShadowImage volume before the second step moves the backup to tape. If you don't, the backup of the previous point-in-time image will be lost.

If you have PSPLIT a number of volumes, and then executed a DUMP with PPRC=, but that DUMP step failed for some reason before processing all of the volumes, you can just resubmit the job. ABR will bypass the backup of any volume for which PSPLIT was done and the backup has already been moved to tape. However, if you have included in the DUMP step other volumes which are to be backed up normally (e.g., volumes not capable of ShadowImage), they will be backed up again by the restarted backup step.

**MULTIPLE
SYSTEMS**

Although the PSPLIT step and the DUMP step can execute on different MVS systems, **FDR InstantBackup for ABR requires that both systems must be able to access both the source and target volumes and a common ABR catalog.** During the PSPLIT, ABR updates the online volume, the offline copy and the ABR catalog. During the DUMP, ABR must read the online volume, and update the offline copy and the ABR catalog.

It is not possible to PSPLIT a ShadowImage volume from one system, and then do the ABR backup on another system (either another LPAR or a remote system) unless that system can address both disk volumes and update the same ABR catalog. However, if you only need full-volume backups (no incremental backups), it is possible to do this outside of ABR with PGM=FDR backups; contact Innovation for guidance.

FDR InstantBackup for ABR does not currently support remote copies of volume using PPRC; the ShadowImage volume must be in the same 7700/7700E subsystem as its online volume.

CONTINUED . . .

27.07 EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK

These are examples of FDR InstantBackup usage with ShadowImage. The examples contain notes on when you should execute various ShadowImage commands, but you should consult current documentation from HdsHitachi Data Systems to verify the proper procedures.

A STEPLIB DD may be required in FDR InstantBackup steps if you have installed the product in a library other than your standard FDR program library.

**FDR FULL
VOLUME
BACKUP**

The ShadowImage volume at address 01FA has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous one-time (C)ESTPAIR or PPRCOPY ESTPAIR command has been issued to create the duplex pair. Step SPLIT will suspend the pair, making the ShadowImage volume accessible to FDR. Step BACKUP will backup the ShadowImage copy of volume PROD01. After the backup is complete, step RESYNC will resynchronize it with its online volume.

In the sample PPRCOPY commands, the "triplet" of numbers in the PRI and SEC operands identify the subsystem ID (SSID), serial number and device address of the primary (online) and secondary (ShadowImage) volumes. You must substitute the values appropriate for your configuration and the disks to be paired. With ShadowImage the serial number will be the same for both but the SSID might be the same or different. You can get this "triplet" of identifying information for any disk device with the TSO command CQUERY DEVN(X'uuuu').

```
//SPLIT EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
PPRCOPY SUSPEND DDNAME(DISK1) PRI(X'0080',30158,X'04') -
SEC(X'0080',30158,X'1A')
//BACKUP EXEC PGM=FDR,REGION=1M2M,COND=(0,NE,SPLIT)
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR.USE.UNIT01FA.DISBAND,UNIT=SYSALLDA,
// VOL=SER=PROD01,DISP=OLD
//TAPE1 DD DSN=BACKUP.VPROD01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSIN DD *
DUMP TYPE=FDR
//RESYNC EXEC PGM=ICKDSF,COND=(0,NE,BACKUP)
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'04') -
SEC(X'0080',30158,X'1A') MODE(RESYNC)
```

The output from such a backup will look something like this (the FDR219 message will confirm that the offline ShadowImage volume was actually backed up):

```
FDR303 CARD IMAGE -- DUMP TYPE=FDR
FDR219 FDR VOL=PROD01 IS BEING DUMPED FROM OFFLINE UNIT=01FA
FDR007 STARTING TIME OF ACTUAL DUMP -- 14.20.27 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR007 ENDING TIME OF ACTUAL DUMP -- 14.20.38 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR122 BYTES DSK TRK T BLKS RESTART STIMERS ERRS ACT DSK LOW HGH DEXCP NUMDS
FDR122N 0022053561 000618 000491 0000000 0000000 000 000618 000 000 00100 00010
FDR002 FDR DUMP SUCCESSFULLY COMPLETED FROM VOL=PROD01
FDR999 FDR SUCCESSFULLY COMPLETED
```

27.07 CONTINUED

**FDR FULL
VOLUME
BACKUP
(CONTINUED)**

The backup created by this example is a normal FDR full-volume backup. There are no special considerations for restoring the entire disk volume or individual data sets from it. For example:

```
//RESTORE EXEC PGM=FDR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPE1 DD DSN=BACKUP.VPROD01(0),DISP=OLD
//DISK1 DD UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
//SYSIN DD *
RESTORE TYPE=FDR
```

**FDRDSF DATA
SET BACKUP**

The ShadowImage volumes at addresses 01F2 and 01F3 are part of a pool. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step SYNC will create the duplex pairs and copy all of the data to the ShadowImage volumes; allow time for the synchronization to complete. At that point updates to the online volumes can be quiesced and the rest of the steps executed. Step SPLIT will detach (reset to simplex) the ShadowImage volume; when it is complete, updates to the online volumes can resume. Step BACKUP will backup the ShadowImage copies of selected data sets on volumes CICS01 and CICS02, processing the two backups in parallel.

```
//SYNC EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=CICS01,DISP=OLD
//DISK2 DD UNIT=SYSALLDA,VOL=SER=CICS02,DISP=OLD
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'13') -
SEC(X'0080',30158,X'32')
PPRCOPY ESTPAIR DDNAME(DISK2) PRI(X'0080',30158,X'17') -
SEC(X'0080',30158,X'33')
//SPLIT EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=CICS01,DISP=OLD
//DISK2 DD UNIT=SYSALLDA,VOL=SER=CICS02,DISP=OLD
PPRCOPY DELPAIR DDNAME(DISK1) PRI(X'0080',30158,X'13') -
SEC(X'0080',30158,X'32')
PPRCOPY DELPAIR DDNAME(DISK2) PRI(X'0080',30158,X'17') -
SEC(X'0080',30158,X'33')
//BACKUP EXEC PGM=FDRDSF,REGION=2M,COND=(0,NE,SPLIT)
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR.USE.UNIT01F2.DISBAND,UNIT=SYSALLDA,
// VOL=SER=CICS01,DISP=OLD
//TAPE1 DD DSN=BACKUP.VCICS01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN1 DD SYSOUT=*
//DISK2 DD DSN=FDR.USE.UNIT01F3.DISBAND,UNIT=SYSALLDA,
// VOL=SER=CICS02,DISP=OLD
//TAPE2 DD DSN=BACKUP.VCICS02(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN2 DD SYSOUT=*
//SYSIN DD *
DUMP TYPE=DSF,ATTACH
SELECT DSN=CICS**
```


27.07 CONTINUED

**FDRCOPY OF
SMS-
MANAGED
DATA**

An SMS storage group consists of 2 volumes, SMS001 and SMS002, which have permanently assigned ShadowImage volumes at addresses 01F4 and 01FA. You want to make a copy of data sets on those volumes but you can only quiesce updates for a few moments. Step DEFINE will define the catalog entries which will allow FDR InstantBackup to access the ShadowImage volumes (you can omit the DEFINE step if you have permanently cataloged the FDR.USE.UNITuuuu.DISBAND data sets). Step SPLIT will suspend the ShadowImage volumes from their online volumes. Step COPY will copy the selected data sets from the ShadowImage volumes, renaming the data sets during the copy; the data sets will be copied to online volumes chosen by SMS as described in [Section 21](#). These copies might be used for parallel batch processing. Step RESYNC will resynchronize them with their online volumes.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE FDR.USE.UNIT01F4.DISBAND NOSCRATCH
DELETE FDR.USE.UNIT01FA.DISBAND NOSCRATCH
SET MAXCC=0
DEFINE NONVSAM(NAME(FDR.USE.UNIT01F4.DISBAND) -
DEVICETYPE(3390) VOLUME(SMS001))
DEFINE NONVSAM(NAME(FDR.USE.UNIT01FA.DISBAND) -
DEVICETYPE(3390) VOLUME(SMS002))

//SPLIT EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=SMS001,DISP=OLD
//DISK2 DD UNIT=SYSALLDA,VOL=SER=SMS002,DISP=OLD
PPRCOPY SUSPEND DDNAME(DISK1) PRI(X'0080',30158,X'03') -
SEC(X'0080',30158,X'14')
PPRCOPY SUSPEND DDNAME(DISK2) PRI(X'0080',30158,X'04') -
SEC(X'0080',30158,X'1A')
//COPY EXEC PGM=FDRCOPY,REGION=0M,COND=(0,NE,SPLIT)
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR.USE.UNIT01F4.DISBAND,DISP=OLD
//DISK2 DD DSN=FDR.USE.UNIT01FA.DISBAND,DISP=OLD
COPY TYPE=DSF
SELECT DSN=PROD.PAYROLL.*,NEWI=..+COPY,STORCLAS=BATCH
//RESYNC EXEC PGM=ICKDSF
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=SYSALLDA,VOL=SER=SMS001,DISP=OLD
//DISK2 DD UNIT=SYSALLDA,VOL=SER=SMS002,DISP=OLD
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'03') -
SEC(X'0080',30158,X'14') MODE(RESYNC)
PPRCOPY ESTPAIR DDNAME(DISK2) PRI(X'0080',30158,X'04') -
SEC(X'0080',30158,X'1A') MODE(RESYNC)
```

27.08 EXAMPLES FOR ABR

**ABR FULL
VOLUME
BACKUP**

ShadowImage volumes have been permanently assigned to the online payroll volumes PAYxxx; a previous one-time ESTPAIR has been issued to establish those pairings.

Step FULL1 will split each ShadowImage volume from its online volume to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
              PSPLIT TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
              MOUNT  VOLG=PAY*
```

The SYSPRINT output of step FULL1 will look something like:

```
FDR304  ABR PSPLIT   REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07C8 ,DSNAME=FDRABR.VPAY001.C1000300 ,FILE=0000,VOL=SER=PAY001
FDR306  PSPLIT      SUCCESSFULLY COMPLETED
FDR304  ABR PSPLIT   REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07C9 ,DSNAME=FDRABR.VPAY002.C1000300 ,FILE=0000,VOL=SER=PAY002
FDR306  PSPLIT      SUCCESSFULLY COMPLETED
FDR304  ABR PSPLIT   REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305  TO OFFLINE UNIT=07CA ,DSNAME=FDRABR.VPAY003.C1000300 ,FILE=0000,VOL=SER=PAY003
FDR306  PSPLIT      SUCCESSFULLY COMPLETED
```

The FDR304 message identifies this as a PSPLIT, and the FDR305 message identifies the ShadowImage target unit address (ABR queries the 7700/7700E to get the address of the ShadowImage volume currently associated with each online volume). It also shows the ABR backup data set name created, and that it was cataloged to the input volume.

The SYSPRINx output will also contain a message to indicate that a PSPLIT was done:

```
FDR218   FDRPPRC UNIT=07C0 IS SPLIT FROM PAIR UNIT=07C8
```

Performance note: in the PSPLIT step, you can improve performance by adding multiple TAPEX DD statements, up to 9, to allow the PSPLITs to operate concurrently. You can put more TAPEX DDs in the PSPLIT step than you have in the actual DUMP step.

Step FULL2 will do the actual backups. PPRC=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the Hitachi Data Systems hardware to get the addresses of the ShadowImage volumes; there is no need to specify them here.

27.08 CONTINUED

```

ABR FULL //FULL2 EXEC PGM=FDRABR,REGION=0M
VOLUME //SYSPRINT DD SYSOUT=*
BACKUP //SYSPRINT1 DD SYSOUT=*
(CONTINUED) //SYSPRINT2 DD SYSOUT=*
//TAPE1 DD DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
// LABEL=EXPDT=99000
//TAPE11 DD DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2 DD DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
// LABEL=EXPDT=99000
//TAPE22 DD DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN DD *
DUMP TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,PPRC=(USE,RET)
MOUNT VOL=PAY*

```

The SYSPRINT output for step FULL2 will be similar to:

```

FDR219 FDR VOL=PAY001 IS BEING DUMPED FROM OFFLINE UNIT=07C8
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY001.C1000300,FILE=0005,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY001.C2000300,FILE=0001,VOL=SER=B15277
FDR218 FDRPPRC UNIT=07C2 WAS REESTABLISHED - MODE(RESYNC)
FDR232 FDRINST DISBANDED PATHS FROM OFFLINE UNIT=07C8
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY002 IS BEING DUMPED FROM OFFLINE UNIT=07C9
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY002.C1000300,FILE=0012,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY002.C2000300,FILE=0001,VOL=SER=B15600
FDR218 FDRPPRC UNIT=07C3 WAS REESTABLISHED - MODE(RESYNC)
FDR232 FDRINST DISBANDED PATHS FROM OFFLINE UNIT=07C9
FDR219 FDR VOL=PAY003 IS BEING DUMPED FROM OFFLINE UNIT=07CA
FDR304 FDR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY003.C1000300,FILE=0006,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY003.C2000300,FILE=0002,VOL=SER=B15277
FDR218 FDRPPRC UNIT=07C4 WAS REESTABLISHED - MODE(RESYNC)
FDR232 FDRINST DISBANDED PATHS FROM OFFLINE UNIT=07CA
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR224 FDR VOL=PAY004 IS BEING DUMPED FROM ONLINE UNIT BECAUSE SNAP/SPLIT WAS NOT ATTEMPTED
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY004,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY004.C1000300,FILE=0013,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY004.C2000300,FILE=0002,VOL=SER=B15600
FDR306 DUMP SUCCESSFULLY COMPLETED

```

FDR219 messages indicate that a previously created point-in-time image was found for the volume being dumped. FDR305 messages show the actual location of the backup data sets. FDR218 messages indicate that the RET option caused ABR to RE-ESTABLISH the ShadowImage volume as an active mirror of its online volume, ready to do the next backup.

In this example, volume PAY004 has no ShadowImage volume. The FDR224 message indicates that there is no point-in-time image of the volume so that a normal ABR backup was done. If you intended to do point-in-time backups all of the volumes in this backup, you must do a TSO CESTPAIR or ICKDSF PPRCOPY ESTPAIR for PAY004 and insure that the FULL1 step includes it; otherwise the backup from PAY004 will not be taken at the same point-in-time as the other volumes. If you have no need to do point-in-time backup on PAY004, you can ignore the message. So, you can include in this ABR step volumes which have no point-in-time requirements; they will be backed up normally.

Note that if you attempted to PSPLIT a volume but the PSPLIT failed for some reason, ABR will set a flag in the ABR Model DSCB on the online volume. The DUMP step will fail the backup of that volume with message FDR211 and a non-zero return code, assuming that you intended it to be a point-in-time backup.

27.08 CONTINUED

**ABR
INCREMENTAL
BACKUP**

The ShadowImage volumes at addresses 01F2 and 01F3 are part of a pool of ShadowImage volumes. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step SYNC may be executed separately before the backup is required, giving the Hitachi Data Systems microcode time to synchronize the ShadowImage volume with the online volume; the step will terminate when the synchronization is complete. At that point updates to the online volumes can be quiesced and the rest of the steps executed.

In the sample PPRCOPY commands, in PRI, X'0080' defines the subsystem ID (SSID) of the primary volume (PROD01), 30158 is the serial number of its disk subsystem, and X'04' is the channel address within the SSID of that disk. SEC provides the same information about the secondary (target) disk. You must substitute the values appropriate for your configuration and the disks to be paired. With ShadowImage the serial number will be the same for both but the SSID might be the same or different. You can get this "triplet" of identifying information for any disk address with the TSO command CQUERY DEVN(X'uuuu').

```
//SYNC      EXEC      PGM=ICKDSF
//SYSPRINT DD      SYSOUT=*
//DISK1     DD      UNIT=SYSALLDA,VOL=SER=CICS01,DISP=OLD
//DISK2     DD      UNIT=SYSALLDA,VOL=SER=CICS02,DISP=OLD
PPRCOPY ESTPAIR DDNAME(DISK1) PRI(X'0080',30158,X'13') -
        SEC(X'0080',30158,X'32')
PPRCOPY ESTPAIR DDNAME(DISK2) PRI(X'0080',30158,X'17') -
        SEC(X'0080',30158,X'33')
```

Step INCR1 will split the ShadowImage volume from its online volume to create a point-in-time image and mark the backup as complete. It will create a new ABR backup cycle and update the online volume with information about the new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that the RETPD= parameter must be specified the same as in the INCR1 step which follows.

```
//INCR1     EXEC      PGM=FDRABR,REGION=0M
//SYSPRINT DD      SYSOUT=*
//SYSPRIN1 DD      SYSOUT=*
//SYSPRIN2 DD      SYSOUT=*
//TAPE1     DD      DUMMY
//SYSIN     DD      *
PSPLIT TYPE=ABR,DSNENQ=USE,RETPD=14,ENQERR=NO
MOUNT VOL=CICS01,PPRCUNIT=01F2
MOUNT VOL=CICS02,PPRCUNIT=01F3
```

Step INCR2 will do the actual backups. PPRC=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the 7700/7700E hardware to get the addresses of the ShadowImage volumes; there is no need to specify them here.

```
//INCR2     EXEC      PGM=FDRABR,REGION=0M
//SYSPRINT DD      SYSOUT=*
//SYSPRIN1 DD      SYSOUT=*
//SYSPRIN2 DD      SYSOUT=*
//TAPE1     DD      DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP)
//SYSIN     DD      *
DUMP TYPE=ABR,DSNENQ=USE,DATA=USED,RETPD=14,EN-
QERR=NO,PPRC=(USE,RET)
MOUNT VOL=CICS01
MOUNT VOL=CICS02
```

For examples of messages and other notes, see the preceding example, "ABR FULL VOLUME BACKUP"

27.08 CONTINUED

ABR RESTORE If it is necessary to do a full-volume restore of one of the volumes involved in the point-in-time backup before the backup is moved to tape, ABR will do it automatically. No special parameters or JCL is required. ABR will restore from the point-in-time image or from tape automatically.

For example, this will restore the four volumes in the previous full-volume backup example:

```
//REST      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT  DD     SYSOUT=*
//SYSIN     DD     *
  RESTORE  TYPE=FDR,DYNTAPE,CONFMESS=NO,ONLINE,CPYVOLID=YES
  SELECT   VOL=PAY001
  SELECT   VOL=PAY002
  SELECT   VOL=PAY003
  SELECT   VOL=PAY004
```

Volumes PAY001-003 will be restored from the point-in-time images on the ShadowImage volumes unless they have been moved to tape. Volume PAY004, since it has no point-in-time image, will always be restored from tape.

However, there are some considerations:

- the point-in-time volume image is always a full-volume image, even if a incremental backup was done. If ABR discovers that the first ABR backup cycle to be restored is a point-in-time image, it simply uses the appropriate hardware features to quickly copy the entire volume back to the target online volume and does not need to read any further backups from tape. Since full-volume restores typically start with the most recent backup of the volume, this provides the fastest recovery.
- but if the most recent ABR backup cycle is on tape, but another lower-numbered cycle is a point-in-time image, ABR will bypass restoring from that image (with a message); any data sets whose most recent backup was on that image may be corrupted or back-level. This could occur if you did a PSPLIT to create a cycle but never did a DUMP to move it to tape, and followed that with a regular ABR backup direct to tape. This is not normal and should not occur if your FDR InstantBackup procedures are setup correctly.
- If you are restoring to a volume other than the original volume (NVOL= on the SELECT statement) and the target volume is not in the same disk subsystem as the point-in-time image, ABR will do a normal disk-to-disk copy to copy all tracks in the image to that volume.
- If ABR attempts to restore from a point-in-time image (ABR backup data set cataloged to the original volume), but that image no longer exists, ABR will bypass that backup. If it was an incremental backup, it will restore the rest of the cycles in that generation. If it was a full-volume backup (cycle 00), the restore will fail (you can uncatalog the backup or specify GEN= and CYCLE= to restore from an earlier generation). This can occur if multiple PSPLIT operations were done without an intervening DUMP to move the backup to tape. It can also occur at a disaster site if the restored ABR catalog is not consistent with the tapes that were brought off-site (some backup entries still indicate the point-in-time backup).

28.01 FDR INSTANTBACKUP FOR AMDAHL TDMF**FDR
INSTANTBACKUP
AND AMDAHL
TDMF**

FDR InstantBackup is an additional cost FDR facility for non-disruptive backup. It works in concert with TDMF, an Amdahl software product which allows for the creation of Point-In-Time copies of existing DASD volumes.

FDR InstantBackup enhances:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY data set copy

It provides non-disruptive backup of offline MVS DASD volumes, both SMS and non-SMS, using technology developed by Innovation Data Processing.

TDMF

Traditional backups may be disruptive of normal operations. It obviously takes some time to backup a data set or volume. If data sets are being updated by some application while they are being backed up, the backup may not be valid and the data sets may not be usable when they are restored. The usual response to this problem is to quiesce all updates to the data during the backup, disrupting normal operation until the backup is complete.

When used with TDMF's Point-In-Time feature, FDR InstantBackup allows you to create a copy of an online disk volume and easily backup or copy data from that frozen copy as if it was coming from the original volume. Updates to the data need to be quiesced only for a few moments to create the frozen copy.

Amdahl's Transparent Data Migration Facility (TDMF) is a hardware-vendor-independent software product used to move DASD volumes from one online device to another non-disruptively. A feature of this product is the ability to create Point-In-Time copies of existing MVS DASD volumes by "freezing" an image of a volume at a particular point in time.

**FDR
INSTANTBACKUP
ENHANCES FDR
COMPONENTS**

FDR InstantBackup enhances FDR, FDRDSF, and FDRCOPY to read from an offline copy of an online volume which was created as a point-in-time frozen image using TDMF. This creates a backup or copy of the volumes or data sets that looks exactly like a normal backup or copy, except that the data is frozen in time.

FDR InstantBackup for TDMF does not currently support FDR/ABR.

28.02 TDMF OPERATION**TDMF
OPERATION**

TDMF Point-in-Time migrations allow the pairing of online and target volumes in order to create an exact copy of the online volume while maintaining physical data integrity. The TDMF "prompt" option should be used so that the operator or user can terminate the pairing and release the target volume when backups are to be run. The TDMF "group" option is recommended when related data exists on multiple volumes which must be backed up at the same point-in-time. Please consult TDMF literature or your Amdahl representative for more information on the operation and features of TDMF.

All data on the online volume is copied the target volume, including the VTOC, VTOCIX, and VVDS. The target volume will retain it's original volume serial number in order to avoid duplicate volume serial numbers at IPL time.

**TDMF WITHOUT
FDR
INSTANTBACKUP**

Although TDMF can create a point-in-time image of a production disk volume, traditional backup products cannot directly use that volume image without additional considerations. The TDMF manual even says that the copy cannot be used on the same CPU as the original volume, but with FDR InstantBackup this is no longer true.

For example, a conventional backup product will require that the TDMF target volume (copy of the original online volume) be brought online. While the target volume does not have the same volume serial number as the online volume, the data sets on the target volume will appear to be uncataloged, which may cause problems during backup and restore unless they are renamed and recataloged.

A conventional backup of the TDMF target volume will appear to be a backup from the new volume serial, not the original volume serial, so additional procedures are required to document this correspondence (e.g., backup of volume B is really a backup of volume A). Restore procedures must be modified to restore the backups to the correct serial (e.g., restore volume B back to volume A and re-label it as A).

There are no special considerations for restores from a backup created by FDR InstantBackup. You can restore entire volumes or individual data sets from those backups. The target volumes for restore will be online volumes, not the TDMF volumes. Since the backups created by FDR InstantBackup appear to be backups of the online volume, there are no special volume serial or data set name concerns.

When FDR InstantBackup is used with FDRCOPY, you can freeze a TDMF volume image of the online volumes containing the desired data sets, and then create copies of those data sets on other online volumes. Data sets can be renamed and cataloged during the copy.

28.03 THE FDR INSTANTBACKUP SOLUTION

**THE FDR
INSTANTBACKUP
SOLUTION**

FDR InstantBackup and TDMF allow you to create and use a duplicate TDMF target volume image *without* bringing it online. FDR full volume backup, FDRDSF logical data set backup and FDRCOPY logical data set copy use FDR InstantBackup technology to read the TDMF target volume even while it remains offline. However, TDMF target volumes retain their original volume serial, so the user must use the offline REFORMAT function of ICKDSF to change the volume serial of the target back to the online volume's serial. A non-disruptive backup or copy can start as soon as the REFORMAT is done. Updates to the online volumes continue while you create the point-in-time backup or copy of the data.

If you are licensed for FDR InstantBackup, you invoke it with FDR JCL which points to the online volume but which contains a special data set name that identifies the offline TDMF volume just split from the online volume. FDR InstantBackup verifies that the TDMF volume is offline and that its volume serial matches the serial of the online online disk. It then will read the offline TDMF volume in exactly the same way that it would read the online disk. FDR InstantBackup produces a backup that appears to be a conventional backup of the online volume serial. FDR InstantBackup blends its unique offline, high-speed non-disruptive backup together with a traditional restore complete with FDR's powerful logical file capability.

28.04 SETTING UP TDMF

To use TDMF, you must install and configure the TDMF software. There are no additional configuration considerations.

**TDMF WITH FDR
INSTANTBACKUP**

1. Execute a TDMF Point-in-Time migration session for any number of volume pairs, using the "prompt" option. This should be done far enough ahead of the scheduled backups so that the volumes will be synchronized by that time.
2. Once the prompt has been issued (volumes are synchronized) and you are ready to do the backup or copy, all updates to the online volumes should be shut down for a brief period of time. Respond to the TDMF prompt, which requests final volume synchronization. After the TDMF job is complete, updates can be restarted. Note that TDMF will mark the target volume offline to MVS.
3. Execute an ICKDSF REFORMAT batch step to rename the offline target volume to the original online volume serial. Do not vary the volume online as MVS will not allow duplicate volume serial numbers.
4. Execute FDR or FDRDSF with the special DISKx DD statement.
5. REFORMAT the offline target volume again in order to avoid duplicate volume serial numbers at IPL time. Alternatively, the target volume may be marked in the IOCP to come up offline.
6. In the event that an IPL occurs for an LPAR that is connected to both the online and target volumes during the FDR InstantBackup process, duplicate volume serial numbers will be detected during MVS startup unless the TDMF target devices are marked offline in the IOCP. The customer must reply properly to ensure that the target volume is taken offline.

**SETTING UP A
TDMF
ENVIRONMENT**

There are 2 ways to setup TDMF use:

- 1) you can permanently assign TDMF target volumes to every online volume for which you plan to use FDR InstantBackup backup or copy. This will require more disk capacity and more device addresses, but is much simpler to administer. Each TDMF target volume is actively mirroring it's online volume during the actual TDMF job execution. Once the prompt is replied to, the TDMF job will terminate and the target volumes will contain the point-in-time image, at which time the backup or copy may be performed. You may want to run TDMF Point-in-time Migration on those volumes all the time, except during the backups, so that you will be ready to do backups at any time.
- 2) you can create a pool of TDMF target volumes and assign them to online volumes as necessary. This requires less disk capacity and fewer device addresses, but is more complex to administer. You will have to activate the online/target pair every time it is to be used. It may take additional time to copy all of the data from the online volume, especially if multiple pairs are being activated. You will require operational procedures to insure that two FDR InstantBackup jobs don't try to use the same TDMF volume for two different online volumes at the same time. You will reset the pair to simplex mode before the backup/copy.

You can use a combination of the two, some online volumes with permanent TDMF volumes (perhaps for volumes with frequent or time-critical backups) and some with TDMF volumes assigned from a pool (perhaps for less frequent or non-time-critical backups).

In either case, the TDMF volumes must match the size (number of cylinders) and format (3380 or 3390) of the online volumes to which they will be assigned.

28.05 FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY

FDR InstantBackup uses standard FDR JCL with one modification. FDR control statements used are standard except for a few considerations. For complete details on other JCL and control statement requirements of FDR, FDRDSF, and FDRCOPY, please see [Sections 10, 20, and 21](#) of the FDR manual, respectively.

FDR InstantBackup uses a special convention for the JCL statements that identify the disk to be backed up or copied, i.e. DISKx DDs. As shown in the examples below, the DD points to the online volume but includes a special data set name which tells FDR the device address of the TDMF volume to be read in place of the online volume. FDR will verify that the volume serial of the online volume is the same as that of the offline TDMF volume.

These special DISKx DDs can be used with:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY input disks for COPY operations only.

They can also be used for FDR full-volume copies (COPY TYPE=FDR) but in most cases it would be more useful to copy the online volume directly.

FDRCOPY MOVE operations from a point-in-time copy are not supported by FDR InstantBackup, since it deletes the input data set. You should always move the live, online version of a data set.

**NON-SMS
VOLUMES**

For non-SMS volumes, a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,UNIT=3390,
//          VOL=SER=vvvvvv,DISP=OLD
```

The UNIT and VOLSER (vvvvvv) must identify the online disk volume. Note that the FDR.USE.UNITuuuu data set does not really exist anywhere in your system, but the presence of the special data set name on the DD statement is very important. It is this name that directs FDR to access the offline TDMF volume device located at UCB address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address. A leading zero must be added if it is a 3-digit address, e.g., DSN=FDR.USE.UNIT01F3. This special data set name is never opened so it will cause no problems with security.

28.05 CONTINUED**SMS-
MANAGED
VOLUMES**

For SMS-managed volumes a DISKx DD statement identifying the disk to be backed up would look like this:

```
//DISKx DD DSN=FDR.USE.UNITuuuu,DISP=OLD
```

Because MVS will ignore a volume serial in the JCL if it points to a SMS-managed volume, the special FDR.USE.UNITuuuu name must be cataloged. **This data set does not really have to exist on the original volume, but it must be cataloged to that volume!** Just as for a non-SMS volume, its presence on the DD statement is very important. It is this name that directs FDR to access the TDMF volume on the offline device located at device address uuuu instead of the original online volume. "uuuu" must be a 4-digit MVS device address; add a leading zero if necessary.

Here is one technique for creating the catalog entry required for FDR InstantBackup. This step will catalog the special FDR.USE.UNITuuuu name to an SMS managed volume. It could be run immediately before a backup.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE FDR.USE.UNITuuuu NOSCRATCH
DEFINE NONVSAM(NAME(FDR.USE.UNITuuuu) -
DEVICETYPE(33xx) VOLUME(vvvvvv))
```

The DEVICETYPE and VOLUME parameters identify the SMS-managed online volume. The special FDR.USE.UNITuuuu name will direct FDR to access the offline TDMF volume device located at device address uuuu instead of the original online SMS managed online volume.

**FDR CONTROL
STATEMENTS**

The control statements used with FDR InstantBackup are the normal statements documented for FDR ([Section 10](#)), FDRDSF ([Section 20](#)), and FDRCOPY ([Section 21](#)) with these modifications:

- Since FDR InstantBackup is not reading the live data, it is never appropriate to use the ENQ= or DSNENQ= operands to enqueue on the VTOC of the TDMF volume or the data sets on the TDMF volume. To do so would unnecessarily prevent access to the live data on the online volume.
- For FDRCOPY, specify DSNENQ=NONE. Also, the input volumes must be specified by DISKx DD statements (to identify the TDMF volume addresses) so you should not select data sets with CATDSN= or VOL=.
- Do not use the FDRCOPY MOVE statement when the input is a TDMF volume. If you do so, the live data sets on the online volume will be deleted at the end of the move.

**FDR
INSTANTBACKUP
OPERATION**

- 1) Execute a TDMF job or jobs, to start the Point-in-time Migration of each selected standard volume to its TDMF target volume. Be sure to do this far enough ahead of your scheduled backup to allow the volumes to be synchronized.
- 2) When ready to backup or copy data, quiesce system and/or application update activity on the standard volumes, if required, and reply to the TDMF prompt to cause final synchronization of the volumes. TDMF will place them offline. Once TDMF ends, you can re-enable updates on the standard volumes.
- 3) Execute an ICKDSF offline REFORMAT to change each target volume to the volume serial of the standard volume.
- 4) Execute FDR, FDRDSF, or FDRCOPY with the special DISKx DD statements described above to access the TDMF target volume.
- 5) When you no longer need the "frozen" data, execute the ICKDSF offline REFORMAT again to return the target volume to its normal volume serial, and vary it online.

CONTINUED ...

28.05 CONTINUED**MULTIPLE
SYSTEMS**

FDR InstantBackup requires that both the original volume and its offline TDMF copy be accessible from the MVS system where the backup is run. If you execute the copy on one system, but execute the backup on a second MVS system which does NOT have access to the online volume, FDR InstantBackup will fail. There is a circumvention for FDR and FDRDSF, but you must contact Innovation for assistance in implementing it.

CONTINUED . . .

28.06 EXAMPLES FOR FDR/FDRDSF/FDRCOPY/FDRCPK

These are examples of FDR InstantBackup usage with TDMF. The examples contain notes on when you should execute various TDMF commands, but you should consult current documentation from AMDAHL to verify the proper procedures.

A STEPLIB DD may be required in FDR InstantBackup steps if you have installed the product in a library other than your standard FDR program library.

**FDR FULL
VOLUME
BACKUP**

The TDMF volume at address 01FA with volser XROD01 has been permanently assigned to online non-SMS volume "PROD01" at address 01E4; a previous TDMF job has been submitted to create the Point-in-time Migration copy of PROD01 onto XROD01. Before executing the job below, quiesce updates to PROD01 and reply to the TDMF prompt to cause final synchronization of the volumes. Step CLIP1 changes the volume serial of the offline volume to that of the online volume (PROD01). Step BACKUP will backup the TDMF copy of volume PROD01 and step CLIP2 will change the volser back to XROD01.

```

/**** REPLY TO THE TDMF PROMPT HERE
//CLIP1 EXEC PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD SYSOUT=*
          REFORMAT UNITADDRESS(01FA) VERIFY(XROD01) VOLID(PROD01)
//BACKUP EXEC PGM=FDR,REGION=1M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//DISK1 DD DSN=FDR. USE .UNIT01FA,UNIT=SYSALLDA,
//          VOL=SER=PROD01,DISP=OLD
//TAPE1 DD DSN=BACKUP.VPROD01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSIN DD *
          DUMP TYPE=FDR
//CLIP2 EXEC PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD SYSOUT=*
          REFORMAT UNITADDRESS(01FA) VERIFY(PROD01) VOLID(XROD01)

```

The output from such a backup will look something like this (the FDR219 message will confirm that the offline target volume was actually backed up):

```

FDR303 CARD IMAGE -- DUMP TYPE=FDR
FDR219 FDR VOL=PROD01 IS BEING DUMPED FROM OFFLINE UNIT=01FA
FDR007 STARTING TIME OF ACTUAL DUMP -- 14.20.27 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR007 ENDING TIME OF ACTUAL DUMP -- 14.20.38 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR122 BYTES DSK TRK T BLKS RESTART STIMERS ERRS ACT DSK LOW HGH DEXCP NUMDS
FDR122N 0022053561 000618 000491 0000000 0000000 000 000618 000 000 00100 00010
FDR002 FDR DUMP SUCCESSFULLY COMPLETED FROM VOL=PROD01
FDR999 FDR SUCCESSFULLY COMPLETED

```

The backup created by this example is a normal FDR full-volume backup. There are no special considerations for restoring the entire disk volume or individual data sets from it. For example:

```

//RESTORE EXEC PGM=FDR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPE1 DD DSN=BACKUP.VPROD01(0),DISP=OLD
//DISK1 DD UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
//SYSIN DD *
          RESTORE TYPE=FDR

```

28.06 CONTINUED

**FDRDSF DATA
SET BACKUP**

The TDMF volumes at addresses 01F2 (TDMF01) and 01F3 (TDMF02) are part of a pool. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. A TDMF job is submitted to create the Point-in-time Migration copies. Before executing the job below, quiesce updates to CICS01/CICS02 and reply to the TDMF prompt to cause final synchronization of the volumes; when it is complete, updates to the online volumes can resume. Step CLIP1 changes the volume serials of the offline volumes to that of the online volumes. Step BACKUP will backup selected data sets from the TDMF copy of those volumes in parallel and step CLIP2 will change the volsers back to TDMF01/TDMF02 for future use.

```

/****  CREATE THE TDMF PAIRS HERE, ALLOW DATA COPY TO COMPLETE
/****  REPLY TO THE TDMF PROMPT WHEN READY TO TAKE THE BACKUPS.
//CLIP2  EXEC  PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD  SYSOUT=*
        REFORMAT UNITADDRESS(01F2) VERIFY(TDMF01) VOLID(CICS01)
        REFORMAT UNITADDRESS(01F3) VERIFY(TDMF02) VOLID(CICS02)
//BACKUP EXEC  PGM=FDRDSF,REGION=2M
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//DISK1   DD   DSN=FDR.USE.UNIT01F2,UNIT=SYSALLDA,
//          VOL=SER=CICS01,DISP=OLD
//TAPE1   DD   DSN=BACKUP.VCICS01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN1 DD  SYSOUT=*
//DISK2   DD   DSN=FDR.USE.UNIT01F3,UNIT=SYSALLDA,
//          VOL=SER=CICS02,DISP=OLD
//TAPE2   DD   DSN=BACKUP.VCICS02(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN2 DD  SYSOUT=*
//SYSIN   DD   *
        DUMP  TYPE=DSF,ATTACH
        SELECT DSN=CICS**
//CLIP2  EXEC  PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD  SYSOUT=*
        REFORMAT UNITADDRESS(01F2) VERIFY(CICS01) VOLID(TDMF01)
        REFORMAT UNITADDRESS(01F3) VERIFY(CICS02) VOLID(TDMF02)

```

CONTINUED ...

28.06 CONTINUED

**FDRCOPY OF
SMS-
MANAGED
DATA**

An SMS storage group consists of 2 volumes, SMS001 and SMS002, which have permanently assigned TDMF volumes at addresses 01F4 and 01FA. You want to make a copy of data sets on those volumes but you can only quiesce updates for a few moments. Step DEFINE will define the catalog entries which will allow FDR InstantBackup to access the TDMF volumes (you can omit the DEFINE step if you have permanently cataloged the FDR.USE.UNITuuuu data sets). See notes in the preceding examples for information on creating the TDMF target volumes and the REFORMAT steps. Step COPY will copy the selected data sets from the TDMF volumes, renaming the data sets during the copy; the data sets will be copied to online volumes chosen by SMS as described in [Section 21](#). These copies might be used for parallel batch processing.

```
//DEFINE      EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
      DELETE FDR.USE.UNIT01F4 NOSCRATCH
      DELETE FDR.USE.UNIT01FA NOSCRATCH
      DEFINE NONVSAM(NAME(FDR.USE.UNIT01F4) -
        DEVICETYPE(3390) VOLUME(SMS001))
      DEFINE NONVSAM(NAME(FDR.USE.UNIT01FA) -
        DEVICETYPE(3390) VOLUME(SMS002))

//***  CREATE THE TDMF PAIRS HERE, ALLOW DATA COPY TO COMPLETE
//***  REPLY TO THE TDMF PROMPT WHEN READY TO TAKE THE BACKUPS.

//CLIP1      EXEC PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD  SYSOUT=*
      REFORMAT UNITADDRESS(01F4) VERIFY(TDMF01) VOLID(SMS001)
      REFORMAT UNITADDRESS(01FA) VERIFY(TDMF02) VOLID(SMS002)
//COPY       EXEC PGM=FDRCOPY,REGION=0M
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//DISK1      DD DSN=FDR.USE.UNIT01F4,DISP=OLD
//DISK2      DD DSN=FDR.USE.UNIT01FA,DISP=OLD
      COPY TYPE=DSF
      SELECT DSN=PROD.PAYROLL.**,NEWI=..+COPY,STORCLAS=BATCH
//CLIP2      EXEC PGM=ICKDSF,PARM=NOREPLYU
//SYSPRINT DD  SYSOUT=*
      REFORMAT UNITADDRESS(01F4) VERIFY(SMS001) VOLID(TDMF01)
      REFORMAT UNITADDRESS(01FA) VERIFY(SMS002) VOLID(TDMF02)
```

29.01 FDR INSTANTBACKUP FOR IBM FLASHCOPY

**FDR
INSTANTBACK
UP AND IBM
FLASHCOPY**

FDR InstantBackup is an additional cost FDR facility for non-disruptive backup. It works in concert with FlashCopy, a feature of the IBM 2105 ESS (Enterprise Storage Server, sometimes called the "Shark") which allows for the "instant" creation of exact duplicates of existing DASD volumes.

FDR InstantBackup enhances:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY data set copy
- FDRABR volume backup (full-volume and incremental backup)

It provides non-disruptive backup of offline MVS DASD volumes, both SMS and non-SMS, using technology developed by Innovation Data Processing.

FLASHCOPY

Traditional backups may be disruptive of normal operations. It obviously takes time to backup a data set or volume. If data sets are being updated by some application while they are being backed up, the backup may not be valid and the data sets may not be usable when they are restored. The usual response to this problem is to quiesce all updates to the data during the backup, disrupting normal operation until the backup is complete.

When used with the FlashCopy feature of the ESS, FDR InstantBackup allows you to instantly freeze a copy of an online disk volume and easily backup or copy data from that frozen copy as if it was coming from the original volume. Updates to the data need to be quiesced only for the few moments necessary to create the frozen copy.

**FDR
INSTANTBACKUP
ENHANCES FDR
COMPONENTS**

FDR InstantBackup enhances FDR, FDRDSF, and FDRCOPY to read from an offline copy of an online volume which was created as a point-in-time frozen image using FlashCopy. This creates a backup or copy of the volumes or data sets that looks exactly like a normal backup or copy, except that the data is frozen in time.

If you are licensed for FDRABR, FDR InstantBackup enhances ABR Volume Backups, both full-volume (TYPE=FDR) and incremental (TYPE=ABR/DSF/AUTO) to use FlashCopy, by creating "instant" frozen copies of the volumes to be backed up, and then moving that captured data to tape.

Note: the IBM FlashCopy implementation does not currently provide any facilities that can be used for data set copies, so the FDR InstantBackup support for FlashCopy does not enhance FDRCOPY for normal disk-to-disk copies.

29.02 FLASHCOPY OPERATION

FLASHCOPY FlashCopy is an additional cost feature of the IBM 2105 ESS, which allows for the "instant" duplication of data sets or entire disk volumes, from one volume to another volume in the same logical subsystem within the ESS. The duplication is not truly instant, since data must actually be copied from one disk to the other, but the internal mechanisms of FlashCopy make it look as if it did take place instantly.

When FlashCopy is invoked, a FlashCopy session is started between the two disks involved. If the session is started with the FlashCopy COPY option, the ESS actually begins copying tracks from the source disk to the target disk. FlashCopy also supports a NOCOPY option, where the tracks will not be physically copied unless necessary.

With either type of FlashCopy session, if a track is updated on the source volume and it has not yet been copied to the target, the contents of the original track, as it existed at the time that the FlashCopy session was started, is immediately copied to the target volume before the update is allowed to proceed. This insures that the "point-in-time" image of the original volume is preserved. Tracks that are not updated on the source volume will be copied by a low-priority background task (COPY option) or not copied at all (NOCOPY option).

The NOCOPY option of FlashCopy is used when the copied data is only needed for a short period of time since the copy is lost when the FlashCopy session is withdrawn (terminated). NOCOPY is used mainly by backup products such as FDR InstantBackup.

The COPY option of Flashcopy is used when the copied data will be used on the target volume even after the FlashCopy session is complete. It is supported by FDR InstantBackup but it is neither required nor recommended.

When an application (such as FDR InstantBackup) reads a track from the target volume which has not yet been physically copied from the source volume, the ESS will actually read that track from the source volume. If a track which has been copied to the target volume is requested, it will be read from the target. So, when the target volume is read, the application will read the tracks exactly as they existed on the source volume at the beginning of the FlashCopy session.

If the application updates the target volume, the operation is similar. If the updated track has not yet been physically copied from the source, it is immediately copied before the update takes place on the target volume. If it has been copied, it is simply updated on the target volume.

The result is that FlashCopy logically creates an instant point-in-time copy of one volume onto another volume in the same ESS. Once the FlashCopy session is established, an application such as FDR InstantBackup can access that frozen point-in-time volume image even though the ESS is still copying tracks from the source to the target volume in the background.

IBM provides TSO commands to invoke FlashCopy. These commands are FCESTABL (establish a FlashCopy session), FCWITHDR (terminate a NOCOPY session or prematurely terminate a COPY session) and FCQUERY (query the status of a session). However, these commands are **not used with** FDR InstantBackup. IBM also provides an API to do these same functions; this API **is used** by FDR InstantBackup.

Note: The 2105 ESS can emulate multiple disk control units. These are called logical control units or logical subsystems. The ESS only supports FlashCopy between two disks which are defined to be in the same logical subsystem.

29.02 CONTINUED

**FLASHCOPY
WITHOUT FDR
INSTANT BACKUP**

Although FlashCopy can create a point-in-time image of a production disk volume, traditional backup products cannot directly use the copied image without additional considerations.

For example, remember that a full-volume FlashCopy will copy all of the original volume's data tracks, including the label track containing the volume serial. Since MVS will not let you put two volumes with the same volume serial online, the new volume is forced offline. A conventional backup product will require that the duplicated volume be brought online, so the volume serial of the target volume must be changed.

Although it is possible to preserve the original volume serial of the target volume, or use ICKDSF to re-label the target volume after the FlashCopy session is started, so that it can be put online, but this invokes additional requirements for SMS managed volumes, VSAM clusters and cataloged data sets. The data sets on the source volume will appear to be uncataloged, which may cause problems during backup and restore unless they are renamed and re-cataloged.

A backup of a re-labeled volume will appear to be a backup from the new volume serial, not the original volume serial, so additional procedures are required to document this correspondence (e.g., backup of volume B is really a backup of volume A). Restore procedures must be modified to restore the backups to the correct serial (e.g., restore volume B back to volume A and re-label it as A).

CONTINUED . . .

29.03 THE FDR INSTANTBACKUP SOLUTION

FDR InstantBackup allows you to create a duplicate volume image with FlashCopy without changing its volume serial or bringing it online. FDR full volume backup, FDRDSF data set backup, FDRCOPY data set copy and FDRABR volume backups use FDR InstantBackup technology to read the duplicated volume even while it remains offline. A non-disruptive backup or copy can start as soon as the FlashCopy session is started. Updates to the original volumes can continue while you create the point-in-time backup or copy of the data.

For FDRABR, FDR InstantBackup makes volume backups a 2-step process. The first step executes a FCOPY statement under FDRABR, to start a FlashCopy session for each selected volume to an offline disk in the same ESS logical subsystem at an address that you provide. This takes only seconds per volume. The copied volume is actually a valid backup at that point-in-time; you can do a restore from it if recovery is required. The second step executes a normal FDRABR full-volume or incremental backup of the same volumes with a parameter which tells ABR to read the offline point-in-time image and move it to a backup on tape. There is no need to identify the target device to the ABR backup step; it knows the address of the most recent copy of each online volume.

FDR InstantBackup also makes FDR or FDRDSF backups and FDRCOPY copies a 2-step process. The first step executes a FCOPY statement under program FDRFLASH, to start a FlashCopy session for each selected volume to an offline disk in the same logical subsystem at an address that you provide. When that step ends, the data has been logically copied and you can resume normal operations on the original volume. Then you execute a FDR or FDRDSF backup or a FDRCOPY copy, with a parameter which tells FDR to read the copy instead of the original volume. It then will read the offline copy in exactly the same way that it would read the online disk. FDR InstantBackup produces a backup that appears to be a conventional backup of the original volume serial. FDR InstantBackup blends its unique off-line, high-speed non-disruptive backup together with a traditional restore complete with FDR's powerful logical file capability.

There are no special considerations for restores from a backup created by FDR InstantBackup. You can restore entire volumes or individual data sets from those backups. The target volumes for restore will be normal online volumes. Since the backups created by FDR InstantBackup appear to be backups of the original volume, there are no special volume serial or data set name concerns.

29.04 SETTING UP FLASHCOPY

**FLASHCOPY
CONFIGURA-
TION**

Once you have installed the proper IBM hardware features and software to support FlashCopy, there are no other requirements for the use of FlashCopy.

However, you will want to reserve a set of device addresses in each subsystem for use as FlashCopy target volumes for FDR InstantBackup. Since a FlashCopy destroys any data already on the target volume, these volumes cannot be used for any other type of storage.

The device addresses assigned to the FlashCopy targets must be defined as normal 3380 or 3390 disks in the disk subsystem in your I/O configuration (IOCP or HCD).

When these disks are in use with FDR InstantBackup, using the FCOPY TYPE=FDR statement, FDR will modify the volume label of the copied volume to insure that it cannot be brought online. To avoid the error messages during IPL due to these modified labels, you may want to mark the FlashCopy target addresses as OFFLINE in your I/O configuration, so that MVS will not attempt to access them at IPL time. FDR InstantBackup will be able access them even though they are offline.

The FlashCopy target volumes should not be eligible for PAV (Parallel Access Volumes) since only FDR will access them.

For example, in HCD, the MVS options for a FlashCopy target device will look similar to:

View Device Parameter / Feature Definition				Row 1 of 6
Command ===> _____ Scroll ===> PAGE				
Configuration ID . . .	CURPI099	Common Configuration data		
Device number . . .	01E0	Device type	3390	
Generic / VM device type	3390			
Parameter/ Feature	Value	Req.	Description	
OFFLINE	Yes		Device considered online or offline at IPL	
DYNAMIC	Yes		Device supports dynamic configuration	
LOCANY	Yes		UCB can reside in 31 bit storage	
ALTCTRL	No		Separate physical control unit path	
SHARED	Yes		Device shared with other systems	
SHAREDUP	No		Shared when system physically partitioned	

**FLASHCOPY
WITH FDR
INSTANTBACKUP**

At any time you can take a full-volume FlashCopy of an online volume to one of the FlashCopy target volumes you have reserved, using the FCOPY statement of program FDRFLASH (for FDR, FDRDSF, or FDRCOPY) or FDRABR (for ABR volume backups). The establishment of the FlashCopy session is almost instantaneous.

If you plan to backup multiple volumes and especially if multiple volumes contain related data which must be backed up at the same point-in-time, you should copy all of those volumes at one time.

As soon as the FCOPY step ends, the data has been preserved and you can begin updating the original volumes while the backups are run.

You never need to bring a target volume online or re-label it or re-initialize it with ICKDSF, as long as it is used only for FDR InstantBackup.

29.04 CONTINUED**SETTING UP A
FLASHCOPY
ENVIRON-
MENT**

There are 2 ways to setup FlashCopy use with FDR InstantBackup:

- 1) you can permanently assign a FlashCopy target volume to every original volume for which you plan to use FDR InstantBackup. This will require more device addresses, but is much simpler to administer, since you always know which FlashCopy target will be used for each original volume.
- 2) you can create a pool of FlashCopy targets and use them with FDR InstantBackup as necessary. This requires fewer device addresses, but is more complex to administer. You will require operational procedures to insure that two FDR InstantBackup jobs don't try to use the same target volume for two different original volumes at the same time.

You can use a combination of the two, some online volumes with permanent targets (perhaps for volumes with frequent or time-critical backups) and some with targets assigned from a pool (perhaps for less frequent or non-time-critical backups).

In either case, the target volumes must match the size (number of cylinders) and format (3380 or 3390) of the original volumes to which they will be assigned. The target volume can be larger than the original volume (more cylinders).

CONTINUED ...

29.05 FDR INSTANTBACKUP FOR ABR

When a normal ABR full-volume or incremental backup completes successfully, it is recorded by doing 2 things:

- ABR catalogs the backup data set it just created in the ABR catalog
- ABR updates the VTOC on the volume to record in the ABR Model DSCB the generation and cycle of the backup just created, and updates the current backup information for every selected data set

So, as soon as the backup data set is complete and available, all ABR records point to it and restores can be immediately requested.

**OPERATION
OF FDR
INSTANTBACK
UP FOR ABR**

With FDR InstantBackup for ABR, all of the same information must be recorded on the online volume, to document the backup. However, these changes can't wait until the backup of the offline point-in-time image to tape is complete, since changes may have occurred to the online volume in the meantime (such as updating a data set again, deleting or allocating data sets).

To solve this, when ABR is invoked to capture the point-in-time image, it will update the online volume immediately. The ABR generation and cycle of the next backup is determined and all of the steps shown above are taken, except that the backup is cataloged to the online volume itself (something that would never occur in a normal ABR backup) so this tells ABR that the backup is a point-in-time image. This functions as a true backup, since you can do restores from it.

Later, ABR will process the point-in-time image and create a normal ABR backup from the offline image, exactly like it would do the backup of the online volume if InstantBackup were not involved. It will recatalog the backup data set to the tape just created. At this point, the backup looks exactly like a normal ABR backup of the online volume.

Even before the backup is moved to tape, ABR can restore the entire volume from the point-in-time image.

**EXECUTING
FDR
INSTANTBACK
UP FOR ABR**

To use FDR InstantBackup for ABR, you must divide your existing ABR volume backup jobs into two steps which may be executed in one job or in two separate jobs:

- 1) the first step will capture all the point-in-time images of the selected primary volumes using the ABR statement FCOPY. This will take only a few seconds per volume selected. We recommend that updates to the online volumes be suspended as far as possible during this step. As soon as it finishes, updates to the online volumes can begin again, since the point-in-time captured image of the volume exists. To insure that the copied volume is not accidentally brought online, ABR will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail
- 2) the second step will create the actual backup tapes, reading the captured point-in-time images. It will be identical to the original ABR step, except that ABR operand FCOPY= must be added to invoke this support. This ABR step can also include backups of volumes which do not have captured images; these volumes will be backed up normally (with a warning message). It also optionally releases the captured images when it is done with them.

When you split the ABR step into two, each step should contain the same JCL and ABR control statements as the original, with a few changes. It is vital that all ABR operands be identical in both steps, especially the TYPE= operand, except as noted below (actually, the second step will fail if the TYPE= operand is different). The examples later in this section illustrate the changes required.

In the first step:

- replace the operation name DUMP on the first ABR control statement with FCOPY
- change all TAPEX and TAPEXX DD statements to DUMMY. However, if the original TAPEX/xx DD statements contained RETPD= or EXPDT= operands, include the same operands on

29.05 CONTINUED

the DUMMY DD statements so that ABR can record the actual expiration of the backups. Similarly, if you coded the ABR operand RETPD= on the DUMP statement, include it on the FCOPY statement as well.

- If the original ABR step used DISKxxxx DD statements or mount statements with VOLG= or STORGRP= to identify the volumes to be dumped, you must change the FCOPY step to provide one MOUNT statement for each volume to be dumped, since the MOUNT statement provides the unit address of the target volume for each FlashCopy. DISKxxxx DD statements can be left in place; it will work with or without them.
- Innovation recommends the use of DSNENQ=USE and ENQ=RESERVE so that ABR can properly protect the volume and data sets during the few moments necessary to create the point-in-time image. However, unlike normal ABR operations, where ABR will still backup a data set or volume even if DSNENQ= finds active data sets, FCOPY will not create the point-in-time if it finds data sets in use, which will also cause the second step to fail trying to backup that image. You can override this by specifying the ENQERR=NO operand on the DUMP statement, or by adding statements such as:

```
SELECT DSN=d s n a m e , DSNENQ=NONE
```

for data sets which you expect to be active at the time of the FCOPY; this is supported even for FCOPY TYPE=FDR

- If the backup includes other volumes in other disk subsystems, or volumes for which you simply don't want to do a point-in-time backup, they should be omitted from this step, even though they can be included in the second step

For example,

```
//STEP1 EXEC PGM=FDRABR,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
FCOPY TYPE=FDR
MOUNT VOL=ABC123,FLASHUNIT=03A4
MOUNT VOL=XYZ321,FLASHUNIT=3B76
```

One or more MOUNT statements must be provided to identify each online source volume to be copied (VOL=) and the target device (FLASHUNIT=) to which to copy it. There must be one MOUNT statement for each volume to be copied, and the FLASHUNIT= value must be unique for each source volume. The online volume identified by VOL= and the target volume identified by FLASHUNIT= must be in the same ESS logical subsystem and the target must be offline to MVS. To insure that the target volume is not accidentally brought online, ABR will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail.

FDRABR will initiate a full-volume FlashCopy NOCOPY session for each selected online volume and offline target. If a FlashCopy session already exists between those two volumes, it will be terminated before the new session is started. If a FlashCopy session already exists between the online volume and a different target volume, FDRFLASH will fail and no FlashCopy session will be started.

FDRABR can also initiate a FlashCopy COPY session by adding the operand FCOPY=COPY to the FCOPY statement. This is not recommended; we suggest that you contact Innovation for assistance before using it.

FCOPY supports one additional operand: VERIFYVOLSER=YES will verify that each FLASHUNIT target volume currently has the same volume serial as the corresponding online source volume. This will be true if the last use of the target was for a FCOPY of this source volume, or if you have used ICKDSF to initialize the target volume. You may need to omit VERIFYVOLSER=YES on the first FCOPY since the volser may not match. VERIFYVOLSER=YES is used to verify that you are not accidentally overlaying a target volume used with another source volume.

CONTINUED . . .

29.05 CONTINUED

In the second step:

- add to the DUMP statement one of the operands
FCOPY=USE - tells ABR to read the offline copy of the online volume specified in JCL. ABR remembers the device address of the offline device most recently used as a target for a copy of each online volume. At the end of the backup, the FlashCopy session will be left active
FCOPY=(USE,REL) - same as FCOPY=USE, but at the end of backing up each volume, FDR will issue a request to the subsystem to terminate (withdraw) the FlashCopy session. This is recommended.
- DSLENQ= and ENQ= will be ignored by ABR when dumping point-in-time volume images, since these would prevent access to the online data sets and prevent changes to the online volume's VTOC, not the offline point-in-time image. However, they will be honored if you backup other non-point-in-time volumes in the same step

Since a target volume can hold only one point-in-time image at a time, you must be sure that another FCOPY is not done to the target volume before the second step moves the backup to tape. If you don't, the backup of the previous point-in-time image will be lost.

If you have copied a number of volumes, and then executed a DUMP with FCOPY=, but that DUMP step failed for some reason before processing all of the volumes, you can just resubmit the job. ABR will bypass the backup of any volume for which FCOPY was done and the backup has already been moved to tape. However, if you have included in the DUMP step other volumes which are to be backed up normally (e.g., volumes not capable of FlashCopy), they will be backed up again by the restarted backup step.

**MULTIPLE
SYSTEMS**

Although the FCOPY step and the DUMP step can execute on different MVS systems, FDR InstantBackup for ABR requires that both systems must be able to access both the source and target volumes and a common ABR catalog. During the FCOPY, ABR updates the online volume, the offline copy and the ABR catalog. During the DUMP, ABR must read the online volume, and update the offline copy and the ABR catalog.

It is not possible to FCOPY a volume from one system, and then do the ABR backup on another system (either another LPAR or a remote system) unless that system can address both disk volumes and update the same ABR catalog. However, if you only need full-volume backups (no incremental backups), it is possible to do this outside of ABR with PGM=FDR backups; contact Innovation for guidance.

CONTINUED . . .

29.06 FDR INSTANTBACKUP FOR FDR/FDRDSF/FDRCOPY

To create the FlashCopy copies of the volumes to be processed by FDR, FDRDSF or FDRCOPY, you must first execute program FDRFLASH with JCL similar to:

```
//COPYVOLS EXEC PGM=FDRFLASH
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DUMMY
//SYSIN DD *
FCOPY TYPE=FDR
MOUNT VOL=PROD01,FLASHUNIT=037D
MOUNT VOL=PROD02,FLASHUNIT=1D97
```

In the FCOPY step, one or more MOUNT statements must be provided to identify each online source volume to be copied (VOL=) and the target device (FLASHUNIT=) to which to copy it. There must be one MOUNT statement for each volume to be copied, and the UNIT= value must be unique for each source volume. The online volume identified by VOL= and the target volume identified by FLASHUNIT= must be in the same logical subsystem and the target must be offline to MVS. To insure that the copied volume is not accidentally brought online, FDRFLASH will modify the volume label in a way that FDR can access it but any attempt to bring it online (VARY command or IPL) will fail.

FDRFLASH will initiate a full-volume FlashCopy NOCOPY session for each selected online volume and offline target. If a FlashCopy session already exists between those two volumes, it will be terminated before the new session is started. If a FlashCopy session already exists between the online volume and a *different* target volume, FDRFLASH will fail and no FlashCopy session will be started.

WARNING: if you execute FDRFLASH to establish a FlashCopy session for a volume, then execute FDRFLASH for that volume again before the FDR/FDRDSF backups or FDRCOPY copies from that FlashCopy are complete, the original point-in-time image will be lost. It is your responsibility to see that the backups or copies complete before another FlashCopy is done.

FDRFLASH can also initiate a FlashCopy COPY session by adding the operand FCOPY=COPY to the FCOPY statement. This is not recommended; we suggest that you contact Innovation for assistance before using it.

WARNING: 4-digit device addresses must be used with FLASHUNIT; specify a leading zero if required.

Ideally, you should quiesce any updates to the source volumes during the FCOPY, to insure that you can get a clean, restorable backup of every data set on the volumes. Once the FDRFLASH step ends (usually within a few seconds), the data has been copied and normal access to the source volumes can be resumed.

FDRFLASH does the same security checks as an FDR full-volume backup ([see Section 10.02](#)).

If you like, you can specify the DSNEQ= operand (to check if data sets are active when copied) and/or the ENQ= operand (to ENQ or RESERVE the VTOC during the copy). [See Section 10.04](#) for details.

FCOPY supports one additional operand: VERIFYVOLSER=YES will verify that each FLASHUNIT target volume currently has the same volume serial as the corresponding online source volume. This will be true if the last use of the target was for a FCOPY of this source volume, or if you have used ICKDSF to initialize the target volume. You may need to omit VERIFYVOLSER=YES on the first FCOPY since the volser may not match. VERIFYVOLSER=YES is used to verify that you are not accidentally overlaying a target volume used with another source volume.

29.06 CONTINUED

The actual backups of the offline point-in-time copies can now be taken with FDR InstantBackup using:

- FDR full volume backup
- FDRDSF data set backup
- FDRCOPY input disks for COPY operations only.

You can also use FDR InstantBackup for FDR full-volume copies (COPY TYPE=FDR) but in most cases it would be more useful to copy the online volume directly.

FDRCOPY MOVE operations from a point-in-time copy are not supported by FDR InstantBackup, since it deletes the input data set. You should always move the live, online version of a data set.

FDR InstantBackup uses standard FDR JCL and control statements with several modifications. For complete details on other JCL and control statement requirements of FDR, FDRDSF, and FDRCOPY, please [see Sections 10, 20, and 21](#) of the FDR manual, respectively. DISKx DD statements in FDR InstantBackup steps must specify the original online disk volume, just as if you were backing it up directly.

To invoke FDR InstantBackup you must add the FCOPY= operand to the DUMP or COPY statement. FCOPY= has two forms:

FCOPY=USE - tells FDR InstantBackup to read the offline copy of the online volume specified in JCL. FDR remembers the device address of the offline device most recently used as a target for a FlashCopy of each online volume unless an intervening IPL has occurred. At the end of the backup, the FlashCopy session will be left active so that additional backups may be done from the point-in-time target if desired.

FCOPY=(USE,REL) - same as FCOPY=USE, but at the end of backing up each volume, FDR will issue a request to the subsystem to terminate (withdraw) the FlashCopy session.

Note: FDR can determine the device address of the offline copy only if FDRFLASH was used to create the copy. If you use TSO commands or any other facility to invoke FlashCopy, FDR will not be able to determine that address.

Other considerations for FDR control statements with FDR InstantBackup:

- Since FDR InstantBackup is not reading the live data, it is never appropriate to use the ENQ= or DSNENQ= operands to enqueue on the VTOC of the FlashCopy target or the data sets on the target. To do so would unnecessarily prevent access to the live data on the online volume.
- For FDRCOPY, specify DSNENQ=NONE. If you use CATDSN= to select data sets to be copied, you must make sure that all of the input volumes which will be selected from the catalog have been copied by a preceding FDRFLASH step.
- Do not use the FDRCOPY MOVE statement when the input is a FlashCopy target. If you do so, the live data sets on the online volume will be deleted at the end of the move.

**FDR
INSTANTBACKUP
OPERATION**

- 1) If you do not have a permanently assigned FlashCopy target volume for every original volume, select target volumes from the pool for use with this execution. You must establish procedures to insure that a given target is not used for another FCOPY until the backup is complete.
- 2) When ready to backup or copy data, quiesce system and/or application update activity on the online volumes, if required, and execute FDRFLASH for each original-target pair. Once the FDRFLASH step ends, you can re-enable updates on the original volumes.
- 3) Execute FDR, FDRDSF, or FDRCOPY with FCOPY=USE or FCOPY=(USE,REL) against each of the online volumes. The offline copy will be read.

CONTINUED . . .

29.06 CONTINUED

MULTIPLE SYSTEMS FDR InstantBackup requires that both the original volume and its offline copy be accessible from the MVS system where the backup is run. If you execute the FCOPY on one system, but execute the backup on a second MVS system which does NOT have access to the online volume, FDR InstantBackup will fail. There is a circumvention for FDR and FDRDSF, but you must contact Innovation for assistance in implementing it.

Even if the volumes are accessible from both systems, the DUMP step will fail if it is run on a different system (another CPU or LPAR) from that where the FCOPY was run. If this is required, contact Innovation for assistance.

CONTINUED ...

29.07 EXAMPLES FOR FDR/FDRDSF/FDRCOPY

These are examples of FDR InstantBackup usage with FlashCopy.

A STEPLIB DD may be required in FDR InstantBackup steps if you have installed the product in a library other than your standard FDR program library.

**FDR FULL
VOLUME
BACKUP**

The FlashCopy target volume at address 01FA has been permanently assigned to online volume "PROD01". The step FCOPY will start a FlashCopy session from online volume PROD01 to offline target 01FA. VERIFYVOLSER=YES will verify that 01FA already contains volser PROD01; you may need to omit this the first time that you FCOPY to that target. Step BACKUP will backup the offline copy of volume PROD01. After the backup is complete, FDR will terminate the FlashCopy session.

```
//FCOPY      EXEC      PGM=FDRFLASH
//SYSPRINT DD      SYSOUT=*
//TAPE1      DD      DUMMY
//SYSIN      DD      *
      FCOPY TYPE=FDR,VERIFYVOLSER=YES
      MOUNT VOL=PROD01,FLASHUNIT=01FA
//BACKUP     EXEC      PGM=FDR,REGION=0M,COND=(0,NE,FCOPY)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1      DD      UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
//TAPE1      DD      DSN=BACKUP.VPROD01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSIN      DD      *
      DUMP TYPE=FDR,FCOPY=(USE,REL)
```

The output from the backup step will look something like this (the FDR219 message will confirm that the offline copy was actually backed up):

```
FDR303 CARD IMAGE -- DUMP TYPE=FDR,FCOPY=(USE,REL)
FDR219 FDR VOL=PROD01 IS BEING DUMPED FROM OFFLINE UNIT=01FA
FDR007 STARTING TIME OF ACTUAL DUMP -- 14.20.27 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR007 ENDING TIME OF ACTUAL DUMP -- 14.20.38 -- UNIT=3390-3 ,IN=DISK1 ,OUTPUT=TAPE1
FDR122 BYTES DSK TRK T BLKS RESTART STIMERS ERRS ACT DSK LOW HGH DEXCP NUMDS
FDR122N 0022053561 000618 000491 0000000 0000000 000 000618 000 000 00100 00010
FDR002 FDR DUMP SUCCESSFULLY COMPLETED FROM VOL=PROD01
FDR999 FDR SUCCESSFULLY COMPLETED
```

The backup created by this example is a normal FDR full-volume backup. There are no special considerations for restoring the entire disk volume or individual data sets from it. For example:

```
//RESTORE EXEC      PGM=FDR,REGION=2M
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//TAPE1      DD      DSN=BACKUP.VPROD01(0),DISP=OLD
//DISK1      DD      UNIT=SYSALLDA,VOL=SER=PROD01,DISP=OLD
//SYSIN      DD      *
      RESTORE TYPE=FDR
```

29.07 CONTINUED

**FDRDSF DATA
SET BACKUP**

The FlashCopy target volumes at addresses 01F2 and 01F3 are part of a pool of targets. In this run they will be used to copy volumes CICS01 and CICS02. When ready to execute, updates to the original volumes may need to be quiesced. The step FCOPY will start FlashCopy sessions from the original volumes to the targets; when the step completes, updates to the original volumes can begin again. Step BACKUP will backup the FlashCopy copies of selected data sets on volumes CICS01 and CICS02, processing the two backups in parallel. After the backup is complete, FDR will terminate the FlashCopy sessions.

```
//FCOPY      EXEC      PGM=FDRFLASH
//SYSPRINT DD      SYSOUT=*
//TAPE1      DD      DUMMY
//SYSIN      DD      *
FCOPY TYPE=FDR
MOUNT VOL=CICS01,FLASHUNIT=01F2
MOUNT VOL=CICS02,FLASHUNIT=01F3
//BACKUP     EXEC      PGM=FDRDSF,REGION=2M,COND=(0,NE,FCOPY)
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//DISK1      DD      UNIT=SYSALLDA,VOL=SER=CICS01,DISP=OLD
//TAPE1      DD      DSN=BACKUP.VCICS01(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN1 DD      SYSOUT=*
//DISK2      DD      UNIT=SYSALLDA,VOL=SER=CICS02,DISP=OLD
//TAPE2      DD      DSN=BACKUP.VCICS02(+1),UNIT=TAPE,DISP=(,CATLG)
//SYSPRIN2 DD      SYSOUT=*
//SYSIN      DD      *
DUMP TYPE=DSF,ATTACH,FCOPY=(USE,REL)
SELECT DSN=CICS**
```

The backups created by this example are normal FDRDSF data set backups. There are no special considerations for restoring data sets from them. For example:

```
//RESTORE EXEC      PGM=FDRDSF,REGION=2M
//SYSPRINT DD      SYSOUT=*
//SYSUDUMP DD      SYSOUT=*
//TAPE1      DD      DSN=BACKUP.VCICS01(0),DISP=OLD
//SYSIN      DD      *
RESTORE TYPE=DSF,RECAT,VRECAT
SELECT DSN=CICS3.LOADLIB
```

CONTINUED ...

29.07 CONTINUED

**FDRCOPY
FROM A
POINT-IN-TIME
IMAGE**

An SMS storage group consists of 2 volumes, SMS001 and SMS002, which have permanently assigned targets at addresses 01F4 and 01FA. You want to make a copy of data sets on those volumes but you can only quiesce updates for a few moments. The step FCOPY will start FlashCopy sessions for the online volumes to the offline target volumes. Step COPY will copy the selected data sets from the copied volumes, renaming the data sets during the copy; the data sets will be copied to online volumes chosen by SMS as described in [Section 21](#). These copies might be used for parallel batch processing.

```
//FCOPY      EXEC      PGM=FDRFLASH,REGION=0M
//SYSPRINT   DD        SYSOUT=*
//SYSIN       DD        *
FCOPY TYPE=FDR
MOUNT VOL=SMS001,FLASHUNIT=01F4
MOUNT VOL=SMS012,FLASHUNIT=01FA
//COPY       EXEC      PGM=FDRCOPY,REGION=0M,COND=(0,NE,FCOPY)
//SYSPRINT   DD        SYSOUT=*
//SYSUDUMP   DD        SYSOUT=*
//DISK1      DD        UNIT=SYSALLDA,VOL=SER=SMS001,DISP=OLD
//DISK2      DD        UNIT=SYSALLDA,VOL=SER=SMS002,DISP=OLD
COPY TYPE=DSF,FCOPY=(USE,REL)
SELECT DSN=PROD.PAYROLL.***,NEWI=...+COPY,STORCLAS=BATCH
```

CONTINUED ...

29.08 EXAMPLES FOR ABR

**ABR FULL
VOLUME
BACKUP**

FlashCopy target volumes have been permanently assigned to the online payroll volumes PAYxxx.

Step FULL1 will start FlashCopy sessions for each online volume to its FlashCopy target to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
FCOPY        TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
MOUNT        VOL=PAY001,FLASHUNIT=07C8
MOUNT        VOL=PAY002,FLASHUNIT=07C9
MOUNT        VOL=PAY003,FLASHUNIT=07CA
```

The SYSPRINT output of step FULL1 will look something like:

```
FDR304  ABR FCOPY      REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305    TO OFFLINE  UNIT=07C8 ,DSNAME=FDRABR.VPAY001.C1000300 ,FILE=0000,VOL=SER=PAY001
FDR306  FCOPY        SUCCESSFULLY COMPLETED
FDR304  ABR FCOPY      REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305    TO OFFLINE  UNIT=07C9 ,DSNAME=FDRABR.VPAY002.C1000300 ,FILE=0000,VOL=SER=PAY002
FDR306  FCOPY        SUCCESSFULLY COMPLETED
FDR304  ABR FCOPY      REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305    TO OFFLINE  UNIT=07CA ,DSNAME=FDRABR.VPAY003.C1000300 ,FILE=0000,VOL=SER=PAY003
FDR306  FCOPY        SUCCESSFULLY COMPLETED
```

The FDR304 message identifies this as a FCOPY, and the FDR305 message identifies the FlashCopy target unit. It also shows the ABR backup data set name created, and that it was cataloged to the input volume.

29.08 CONTINUED

Step FULL2 will do the actual backups. FCOPY=(USE,REL) tells ABR to determine if a point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR remembers the unit addresses of the offline volumes specified in the FULL1 step; there is no need to specify them here.

```
//FULL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11     DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2      DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22     DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN      DD     *
              DUMP TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,FCOPY=(USE,REL)
              MOUNT VOL=PAY*
```

The SYSPRINT output for step FULL2 will be similar to:

```
FDR219 FDR VOL=PAY001 IS BEING DUMPED FROM OFFLINE UNIT=07C8
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY001,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY001.C1000300,FILE=0005,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY001.C2000300,FILE=0001,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY002 IS BEING DUMPED FROM OFFLINE UNIT=07C9
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY002,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY002.C1000300,FILE=0012,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY002.C2000300,FILE=0001,VOL=SER=B15600
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR219 FDR VOL=PAY003 IS BEING DUMPED FROM OFFLINE UNIT=07CA
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL1,VOL=SER=PAY003,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE1,DSNAME=FDRABR.VPAY003.C1000300,FILE=0006,VOL=SER=B10023
FDR305 TO TAPE DDNAME=TAPE11,DSNAME=FDRABR.VPAY003.C2000300,FILE=0002,VOL=SER=B15277
FDR218 FDRINST UNIT=07C2 WAS RELEASED
FDR306 DUMP SUCCESSFULLY COMPLETED
FDR224 FDR VOL=PAY004 IS BEING DUMPED FROM ONLINE UNIT BECAUSE SNAP/SPLIT WAS NOT ATTEMPTED
FDR304 ABR DUMP REQUEST FOR DDNAME=DISKONL2,VOL=SER=PAY004,UNIT=3390-3
FDR305 TO TAPE DDNAME=TAPE2,DSNAME=FDRABR.VPAY004.C1000300,FILE=0013,VOL=SER=B10452
FDR305 TO TAPE DDNAME=TAPE22,DSNAME=FDRABR.VPAY004.C2000300,FILE=0002,VOL=SER=B15600
FDR306 DUMP SUCCESSFULLY COMPLETED
```

FDR219 messages indicate that a previously created point-in-time image was found for the volume being dumped. FDR305 messages show the actual location of the backup data sets. FDR218 messages indicate that the REL option caused ABR to terminate the FlashCopy session.

In this example, volume PAY004 was not copied. The FDR224 message indicates that there is no point-in-time image of the volume so that a normal ABR backup will be done. If you intended to do point-in-time backups all of the volumes in this backup, you must do a FCOPY for PAY004 in the FULL1 step; otherwise the backup from PAY004 will not be taken at the same point-in-time as the other volumes. If you have no need to do point-in-time backup on PAY004, you can ignore the message. So, you can include in this ABR step volumes which have no point-in-time requirements; they will be backed up normally.

Note that if you attempted to FCOPY a volume but the FCOPY failed for some reason (perhaps because the volume is not in a FlashCopy-capable subsystem), ABR will set a flag in the ABR Model DSCB on the online volume. The DUMP step will fail the backup of that volume with message FDR211 and a non-zero return code, assuming that you intended it to be a point-in-time backup.

CONTINUED . . .

29.08 CONTINUED

**ABR
INCREMENTAL
BACKUP**

The FlashCopy target volumes at addresses 01F2 and 01F3 are part of a pool of targets. They are temporarily assigned to online volumes CICS01 and CICS02 at addresses 01D3 and 01D7. Step INCR1 will start a FlashCopy session for each online volume to create a point-in-time image and mark the backup as complete. It will create a new ABR backup cycle and update the online volume with information about the new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that the RETPD= parameter must be specified the same as in the INCR1 step which follows.

```
// INCR1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY
//SYSIN       DD     *
FCOPY  TYPE=ABR,DSNENQ=USE,RETPD=14,ENQERR=NO
MOUNT  VOL=CICS01,FLASHUNIT=01F2
MOUNT  VOL=CICS02,FLASHUNIT=01F3
```

Step INCR2 will do the actual backups. FCOPY=(USE,REL) tells ABR to determine if a point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that remembers the addresses of the offline targets; there is no need to specify them here.

```
// INCR2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP)
//SYSIN       DD     *
DUMP  TYPE=ABR,DSNENQ=USE,DATA=USED,RETPD=14,ENQERR=NO,FCOPY=(USE,REL)
MOUNT  VOL=CICS01
MOUNT  VOL=CICS02
```

For examples of messages and other notes, see the preceding example, "ABR FULL VOLUME BACKUP"

CONTINUED ...

29.08 CONTINUED

ABR RESTORE If it is necessary to do a full-volume restore of one of the volumes involved in the point-in-time backup before the backup is moved to tape, ABR will do it automatically. No special parameters or JCL is required. ABR will restore from the point-in-time image or from tape automatically.

For example, this will restore the four volumes in the previous full-volume backup example:

```
//REST      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT  DD     SYSOUT=*
//SYSIN     DD     *
RESTORE TYPE=FDR,DYNTAPE,CONFMESS=NO,ONLINE,CPYVOLID=YES
SELECT VOL=PAY001
SELECT VOL=PAY002
SELECT VOL=PAY003
SELECT VOL=PAY004
```

Volumes PAY001-003 will be restored from the point-in-time images on the target volumes unless they have been moved to tape. Volume PAY004, since it has no point-in-time image, will always be restored from tape.

However, there are some considerations:

- the point-in-time volume image is always a full-volume image, even if a incremental backup was done. If ABR discovers that the first ABR backup cycle to be restored is a point-in-time image, it simply copies the entire volume back to the target online volume and does not need to read any further backups from tape. Since full-volume restores typically start with the most recent backup of the volume, this provides the fastest recovery.
- but if the most recent ABR backup cycle is on tape, but another lower-numbered cycle is a point-in-time image, ABR will bypass restoring from that image (with a message); any data sets whose most recent backup was on that image may be corrupted or back-level. This could occur if you did a FCOPY to create a cycle but never did a DUMP to move it to tape, and followed that with a regular ABR backup direct to tape. This is not normal and should not occur if your FDR InstantBackup procedures are setup correctly.
- If you are restoring to a volume other than the original volume (NVOL= on the SELECT statement) and the target volume is not in the same disk subsystem as the point-in-time image, ABR will do a normal disk-to-disk copy to copy all tracks in the image to that volume.
- If ABR attempts to restore from a point-in-time image (ABR backup data set cataloged to the original volume), but that image no longer exists, ABR will bypass that backup. If it was an incremental backup, it will restore the rest of the cycles in that generation. If it was a full-volume backup (cycle 00), the restore will fail (you can uncatalog the backup or specify GEN= and CYCLE= to restore from an earlier generation). This can occur if multiple FCOPY operations were done without an intervening DUMP to move the backup to tape. It can also occur at a disaster site if the restored ABR catalog is not consistent with the tapes that were brought off-site (some backup entries still indicate the point-in-time backup).

CONTINUED . . .

THIS PAGE WAS INTENTIONALLY LEFT BLANK

30.01 FDRREORG OVERVIEW

FDRREORG is a separately-priced component of the FDR DASD Management System. Its purpose is to reorganize those data sets that require reorganization, namely:

- ICF VSAM KSDS clusters
- IAM (Innovation Access Method) data sets
- PDS (Partitioned) data sets

FDRREORG consists of several parts:

- The FDRREORG program, described in this section, [Section 30](#), automates the reorganization of the data set types listed above.
- The FDRCOPY program, although it is a part of the base FDR product (described in [Section 21](#)), it is enhanced when FDRREORG is licensed to include a REORG statement, used to reorganize PDSs. The REORG statement is described in [Section 32](#). FDRCOPY REORG is invoked internally by FDRREORG to reorganize PDSs, and it can also be executed directly.
- FDRREORG has its own option table, separate from the FDR Global Option Table. The FDRREORG program is used to set and display FDRREORG options and defaults. It is described in [Section 31](#).
- FDRREORG can be invoked for individual data sets from the FDR ISPF panels, as described in [Section 33](#).

**FDRREORG
OVERVIEW**

FDRREORG will logically reorganize IAM and VSAM KSDSs or AIXs, and compress partitioned data sets, based on user specified exclusion and selection criteria. Data sets can be reorganized on an as needed basis using keywords that define the characteristics of data sets requiring reorganization. FDRREORG can also be run in simulation mode to obtain a report of those data sets meeting the users selection criteria.

FDRREORG will attempt to obtain exclusive control of a data set during its reorganization. For IAM data sets, FDRREORG will issue the appropriate IAM ENQs to lock out all updates to the data set while it is being reorganized. For VSAM data sets, FDRREORG will set the update inhibit indicator for the data component of the VSAM cluster being reorganized. If FDRREORG is unable to obtain exclusive control, the data set will be bypassed if the default DSNRETRY option is being used. FDRREORG can be instructed to retry data sets that were in use and can also issue the appropriate system ENQs to process a data set when it becomes available. If it is critical that a particular data set or set of data sets be reorganized, FDRREORG can be instructed to wait for these data sets to become available.

All IAM or VSAM data sets selected for reorganization are backed up to either TAPE or disk, and then immediately reloaded from the backup. The backups are logical backups obtained by using the standard access method interfaces. The striped file format can be used for backups to disk. VSAM data sets with records that are larger than 32763 are not supported. The FDRREORG backups can be used to reload a data set using IDCAMS REPRO or any other standard file copy utility that supports VSAM data sets. If you are creating compressed IAM backups, you must use the BACKUPCOMPRESSED IAM override to use the backup with IDCAMS REPRO. See the IAM manual for additional information. The user has the option of creating the backup data sets as either standard sequential data sets or as GDG's. If GDG's are chosen, FDRREORG will dynamically define the base GDG to an ICF catalog if one does not already exist. GDG's cannot be used with the parallel backup option. The name of the backup data sets are generated dynamically from the target data set name using the BACKUPGROUP or BACKUPINDEX keywords to provide user control of the backup data set name. All backup data sets are cataloged after they have been created and will be uncataloged and deleted after the reload has completed if the backups are not to be kept. FDRREORG provides last tape support on a jobname basis for users who wish to have subsequent executions of the same reorg job append new backups to the last tape used in the last execution of FDRREORG.

30.01 CONTINUED . . .

FDRREORG OVERVIEW (continued) FDRREORG provides complete dynamic allocation of all target and backup data sets. There are numerous keywords provided to allow the user to control the type and number of devices allocated by FDRREORG. Backups on tape will automatically be stacked up to the user specified MAXFILE value.

FDRREORG has a multi-tasking capability that can be activated by providing selection criteria that allows multiple volumes to be processed. Each volume selected for processing will be processed by a separate subtask. Up to 15 concurrent subtasks can be active provided there is sufficient virtual storage available within the region. To minimize the below 16M virtual storage requirements, VSAM buffers and control blocks will be allocated above the 16M. IAM provides this capability automatically for the backup and will provide it for the reload if BSAM mode has not been specified for IAM file load processing via the IAM global option table. QSAM buffers used for the backup will also be allocated above the 16M line if possible.

In order to minimize the manual effort required to recover from a failure during FDRREORG processing, a checkpoint and log file are created to record information that simplifies the recovery of a failed reorganization. The checkpoint file records information on the current status of all active file reorganizations and the log file is used to record information about unsuccessful reloads. FDRREORG's RECOVER statement can dynamically allocate the checkpoint and log files and use the information in these files to complete reload processing for any or all of the data sets impacted by a system or FDRREORG failure.

FDRREORG PERFORMANCE FDRREORG has been designed to provide capability and ease of use in reorganizing VSAM, IAM and PDS data sets without sacrificing performance. FDRREORG can reorganize simultaneously 1 to 15 volumes (default set at 4). FDRREORG gives users intelligent choices in reorganization. These features plus the special functions detailed in the following paragraphs give FDRREORG unrivalled performance and capability.

VSAM PERFORMANCE Normally, VSAM data sets must be deleted and defined before they can be reloaded by IDCAMS. FDRREORG does not need to delete and define most VSAM data sets even if they are marked NOREUSE. VSAM defaults to recovery on a load of a KSDS or AIX. This requires VSAM to preformat each CA prior to loading records into the CA. FDRREORG eliminates the need to reformat the CA's even if recovery is on.

Most important for performance, FDRREORG calculates the optimal buffers required to sequentially read and reload the file.

FDRREORG reduces the wall clock, CPU and I/O operations required to reorganize your VSAM files by 20 to 60%.

PDS PERFORMANCE FDR uses a proprietary technique to compress PDS data sets which reduces the wall clock, CPU and disk EXCPs to compress source libraries by 30 to 60% and load libraries by 80 to 98% as compared to IEBCOPY performing the compression. FDR is capable of compressing PDS's from many volumes without individual JCL and control statements.

SYS1 DATA SETS FDRREORG does not allow you to compress SYS1 PDS data sets directly. These data sets can be compressed with the REORG statement of FDRCOPY which does not look at the FDRREORG list of unmovable data sets ([see section 32.03](#) for FDRCOPY examples).

PARALLEL OPTION Many users have IAM and VSAM files that exceed 1 GB in size. Some IAM files exceed 10GB. Reorganizing large files can take a very long time reducing their availability to online systems. FDRREORG Parallel option backs up each volume of a multi-volume file concurrently to separate tape or disk files (up to a specified maximum). For example, if a file resides on 4 disk volumes the backup time can be reduced by 75%.

3380 TO 3390 While FDRCOPY and FDRDSF can be used to move VSAM KSDS data sets from 3380 to 3390, FDRREORG can also be used to move VSAM KSDS files between device types with the added flexibility of changing the primary and secondary allocation quantities, in addition to changing the CI/CA free space and reorganizing the cluster during the move.

30.01 CONTINUED . . .

**SAMPLE PDS
COMPRESSION
TEST**

User compressed 479 PDS data sets using FDRCOPY REORG as compared to DF/DSS (executing IEBCOPY). Data sets occupied 33585 tracks on a 3380 disk. A mix of source and load libraries were compressed.

	# DSN	ELAPSED	CPU	EXCPS
FDRCOPY	479	10.16M	16.01S	5197
DFDSS V2.5	479	164.39M	382.86S	447203
SAVINGS		93.8%	95.8%	98.8%

**COMPRESS
LINKLIB AND
SMPPTS**

User compressed SYS1.LINKLIB and SYS1.SMPPTS data sets. FDRCOPY was compared to IEBCOPY.

		ELAPSED	CPU	EXCPS
	FDRCOPY	2.52M	7.92S	1136
SYS1.SMPPTS	IEBCOPY	4.50M	23.22S	6074
	SAVINGS	42.8%	65.8%	81.2%

		ELAPSED	CPU	EXCPS
	FDRCOPY	0.64M	2.97S	328
SYS1.LINKLIB	IEBCOPY	7.50M	60.51S	20574
	SAVINGS	91.5%	95.1%	98.4%

REORG BASICS

As shipped, FDRREORG has default criteria that should select those data sets that are in the most need of reorganization. Using the data set names or filters that you provide, FDRREORG will automatically reorganize the following data sets:

1. VSAM KSDSs that have taken control interval or control area splits in 10% or more of the total control intervals or control area's respectively, and have a CA and CI freespace value that is less than 50%.
2. IAM data sets that have used at least 90% of the independent overflow area (compatible format only), or that have at least 10% of the records in the file in the independent overflow area, or require more than 1 megabyte for the overflow index.
3. Partitioned data sets that are at least 90% full.

These defaults can be permanently changed or disabled by using the FDRREORG option change utility FDRREOZO. You can also override these defaults at run time by providing your own values.

Using the appropriate SELECT keywords, you can easily provide your own rules for selecting data sets for reorganization. We highly recommend that you use the SIMULATE statement before running a REORG for the first time, or after making changes to an existing set of selection criteria. SIMULATE will produce a report of the data sets that would have been selected for reorganization had the REORG statement been used. This report can be used to fine tune your selection parameters so only the data sets you want are selected for reorganization. Keep in mind that SIMULATE does not test the availability of the data sets selected. It is possible that any number of the data sets that meet your selection criteria will be in use by other jobs when the REORG statement is run and will not be reorganized. FDRREORG performs a full file reload and must have exclusive control of a data set in order to reorganize it.

30.01 CONTINUED . . .

VSAM KSDSs AND AIXs Due to the nature of the insert strategy used by VSAM, KSDSs need to be reorganized from time to time. When records are inserted into a VSAM KSDS, there may be insufficient free space available in the control interval where the inserted record belongs. When this occurs, VSAM will "split" the control interval to insert the record. If there is insufficient free space in the control area to hold this split control interval, a control area split will also occur. Over time, and without periodic reorganization, these control interval and control area splits will eventually cause a degradation in performance, and in the case of control area splits, cause a waste of valuable disk space.

FDRREORG can be used to logically reorganize these files on an as needed basis. By using either the default selection criteria or your own, **FDRREORG could be run on a daily basis against these files, and would only reorganize them when needed.** This is a much more efficient use of valuable computer resources than simply scheduling a file reorganization on a daily or weekly basis. With FDRREORG's approach, you can be certain that files will be reorganized when they meet your specified criteria, instead of by the best guess method you are probably using today.

FDRREORG supports any VSAM KSDS that has been defined with a maximum record length of 32,763 or less. For VSAM KSDSs that are the base cluster of one or more alternate indexes, FDRREORG will delete and define the base cluster and all associated AIXs and Paths. After the base cluster has been reorganized, FDRREORG will perform a fast build of all the related alternate indexes.

IAM (INNOVATION ACCESS METHOD) IAM, Innovation's proprietary access method, provides a transparent replacement for many VSAM KSDS and ESDS applications. IAM provides improved performance, better space utilization, and built in data compression. FDRREORG completely reorganizes IAM files by placing all records from independent overflow and prime extension into prime data blocks. After reorganization, 100% of independent overflow and prime extension are available for inserts or adds.

By using the appropriate keywords, you can re-size independent overflow and prime extension for files in compatible format. This might be desirable if the IAM file was originally created with a single or dummy record load. IAM treats single or dummy record loads as a special case and dynamically increases the size of independent overflow and prime extension depending on the key of the record loaded. If you completely reorganize an IAM file, retaining the generous amounts of independent overflow and prime extension is in all likelihood inappropriate. Sometimes the generous amounts of independent overflow or prime extension may have been provided when the file was defined because it was not clear how much of these area's would be required. In other cases files may have been defined with insufficient independent overflow or prime extension. FDRREORG allows you to dynamically adjust the size of independent overflow and prime extension for these files. By using the PCTOFKEEPSINGLE keyword for files created with a single record load, and the PCTOFKEEPMULTIPLE keyword for files not created with a single record load, you can specify the percentage of independent overflow to retain during reorganization. To dynamically increase the size of independent overflow, specify a value greater than 100. The PCTPEKEEPSINGLE and PCTPEEKEPMULTIPLE keywords provide the same support for IAM's prime extension. IAM's file load has been enhanced to recognize that FDRREORG is re-loading the IAM file and will accept the temporary overrides to re-size independent overflow and prime extension. If the IAM file is subsequently re-loaded using whatever procedure is normally used to create this file, IAM will revert to the original sizes specified when the IAM file was defined. This ensures that your definition parameters will always be preserved.

IAM files with alternate indexes (AIXs) can also be reorganized. Unlike native VSAM AIXs, IAM AIXs do not have to be rebuilt when the base file is reorganized by FDRREORG.

PARTITIONED DATA SETS When existing members of a partitioned data sets are updated, the partitioned data set access method (BPAM) stores the updated member at the end of the data set. The space occupied by previous versions of these updated members become dead space and can not be re-used unless the data set is compressed with a utility such as IEBCOPY. Over time, the PDS will eventually exhaust all the remaining free space and will have to be compressed before additional members can be added or updated members saved. FDRREORG can dynamically compress these data sets based on the default PDSFULL percentage of 90%, or using a value that you specify. If the PDS

30.01 CONTINUED . . .

**DATA SET
INTEGRITY
DURING
REORGANIZATION**

meets the specified PDSFULL percentage and the PDS has just been compressed, there will be no difference in the amount of available free space after FDRREORG has compressed the data set.

To ensure that the data sets being reorganized can not be updated while a reorganization is taking place, FDRREORG will use techniques that are appropriate for the data sets access method to lock out updates by other jobs. FDRREORG will allocate all data sets using DISP=OLD.

Although this can be changed by using the IAMDISP, PDSDISP or VSAMDISP keywords, **we strongly recommend that you allow FDRREORG to allocate using DISP=OLD.**

For VSAM, FDRREORG will alter the files share options to 1,3 and set the update inhibit flag before starting the backup. If another job has the data set opened for update, FDRREORG will be unable to open the file. If this should occur, FDRREORG will restore the files share options, turn off the update inhibit flag, and bypass the data set. For IAM, FDRREORG will issue IAM's ENQs before opening the data set. If exclusive control is not obtained, FDRREORG will bypass the data set. If you use GRS or CA-MIM, we recommend that all IAM and VSAM ENQs be propagated if the data sets selected for reorganization can be in use on other systems in your complex.

**RECOVERING
FROM A
FAILED
REORGANIZATION**

FDRREORG creates a checkpoint and log file when a REORG statement is executed. If you have multiple REORG statements for a single execution of FDRREORG, the same checkpoint and log file will be used. In the event that one of FDRREORG's subtasks abends, the additional REORG statements will not be executed to preserve the information for the failed subtask. In the case of a failure for a single data set, this information will be kept in FDRREORG's log file. FDRREORG will in this case execute the remaining REORG statements, if any.

The data set names used for the checkpoint and log file contain the name of the job, the date the REORG statement was started, and the time the REORG statement was started. In the event of a subtask abend, FDRREORG abend, or system failure, the checkpoint will be kept and can be used with FDRREORG's RECOVER statement. If none of these conditions occur, the checkpoint file will be deleted by FDRREORG when the job step ends. The log file will be kept if any file re-load failures have occurred. The log will also be kept if a system failure or FDRREORG abend occurs, although it is possible that the log file will be empty in this situation.

To actually recover from a failed REORG, execute FDRREORG with the RECOVER statement and code the JOBNAME keyword specifying the name of the job that failed. FDRREORG's recovery processor will allocate the most recent matching checkpoint and log file for the jobname specified, and re-load any data sets that had started the re-load phase without completing. Data sets that failed during the backup phase will not be recovered because they are still loaded and usable.

**DATA SET
AVAILABILITY
AND THE
DSNRETRY
OPTION**

During the course of a reorganization run, it is possible that data sets selected for reorganization will be in use by other jobs. In many cases, these data sets will become available before FDRREORG completes. By using FDRREORG's DSNRETRY options, you can request that FDRREORG make additional attempts to reorganize these data sets. Depending on the importance you associate with reorganizing a particular data set or group of data sets, there are four options that can be specified to control retry processing.

For those data sets that should only be reorganized if they are available at the time FDRREORG first attempts to reorganize them, use DSNRETRY=NO (the default) on the SELECT statement.

For those data sets that you wish to reorganize if they should become available during the course of execution, use DSNRETRY=RETRY.

To improve FDRREORG's chances of successfully allocating an unavailable data set, use DSNRETRY=ENQ. With this option, FDRREORG will issue an ENQ for the data set. This data set will be the next data set processed if it should become available.

For those data sets where reorganization is critical, use DSNRETRY=WAIT. In addition to issuing the ENQ for these data sets, FDRREORG will wait for them to become available before terminating. If you use the WAIT option, we recommend that you specify the RUNTIME or STOPTIME keyword on the REORG statement, or use an operator STOP (P) command, to prevent FDRREORG from waiting indefinitely.

DSNRETRY is not supported in a JES3 installation; use the default of DSNRETRY=NO.

30.01 CONTINUED . . .

**BACKUP
DATA SETS**

In order for FDRREORG to reorganize VSAM and IAM data sets, it must create a logical backup of these data sets. This backup is used to reload the selected data set. The backup data sets created by FDRREORG are standard sequential data sets and can be used with IDCAMS REPRO. If you are creating compressed IAM backups, you must use the BACKUPCOMPRESSED IAM override to use the backup with IDCAMS REPRO. All of the backup data sets will have a record format of VBS and a logical record length of 32767. The block size will be an optimum block size for the backup device. The name of the backup data set is generated using the BACKUPGROUP or BACKUPINDEX specified on either the REORG or SELECT statement, or from the FDRREORG option table. To avoid having the backup data sets cataloged in your master catalog, be sure to define aliases for any high level indexes that you intend to use for FDRREORG backup data sets.

By default, FDRREORG will delete the backup data set after the target data set has been successfully reorganized. If you would like to keep FDRREORG's backup data sets, code BACKUP=PERM or BACKUP=GDG on either the REORG or SELECT statement, or use the option change utility FDRREOZO to change the default. GDG's cannot be used with the parallel backup option. When PERM is used, the name generated by applying the BACKUPGROUP or BACKUPINDEX is left cataloged. You must provide some method of deleting these data sets if FDRREORG is to be run again against the same target data sets. FDRREORG will not be able to reorganize a VSAM or IAM data set if a data set with the same name as the one generated for the backup data set already exists. The GDG option is probably more useful for retaining FDRREORG's backup data sets. With the GDG option, FDRREORG will use the name generated for the backup data set to define a generation data group (if one does not already exist), and will create the backup data sets as a +1 generation. When defining these generation datagroups, FDRREORG will use NOEMPTY, SCRATCH, and LIMIT(5). If you would prefer to use different parameters, you can use the GDGEMPTY, GDGNOSCRATCH, or GDGLIMIT keywords on the appropriate SELECT statement.

**LAST TAPE
SUPPORT**

If you are using BACKUP=PERM or BACKUP=GDG and are creating your backup data sets on tape, you can direct FDRREORG to append new backups in subsequent runs to the same tapes. To do so, specify LASTAPE on the REORG statement. FDRREORG will create a catalog entry on a jobname and subtask basis that is used in next run of the same job to reuse these tapes. The data set name will be generated using either the default high level prefix in the FDRREORG option table which defaults to FDRREORG, or the prefix specified with the LASTAPEPREFIX keyword, followed by the jobname, TASKnn where nn is the taskid, and finally LASTAPE. The catalog entry will also contain the tape volume and the next file sequence number that will be used for this tape volume. In the event that the file sequence number of the last file created is either 255 or the same as the value specified with the MAXFILE keyword, no catalog entry will be created. This feature is not supported with parallel backups.

30.01 CONTINUED . . .

**ERROR
HANDLING**

FDRREORG has been designed to continue processing if an error occurs while reorganizing an individual data set. FDRREORG makes a distinction between environmental errors, system errors, and fatal internal errors. For a fatal internal error, FDRREORG will allow any in progress reorganizations to complete if possible, and terminate. For environmental and system errors, FDRREORG by default, will allow up to 99 of both types to occur before shutting down. Please keep in mind that this discussion applies to errors encountered while backing up and reloading data sets only. Any abends that occur in FDRREORG's main task will result in an immediate abnormal termination. Similarly, any abends that occur in FDRREORG's volume processor tasks will result in the immediate abnormal termination of that task. The remaining volume processor tasks will continue to run, although resources allocated to the failed task may not have been released.

To allow you to control how FDRREORG handles these errors, a number of keywords have been provided. For a volume processor subtask abend, FDRREORG by default will attempt to complete the active REORG function without the abended task and terminate. If there are additional REORG statements, they will not be processed. In addition, any data sets which would have been selected for reorganization on the volume being processed by the abended task will not be selected by the remaining tasks. If you would prefer that FDRREORG shutdown when a volume processor subtask abends, specify SUBTASKABEND=TERM on the REORG statement. FDRREORG will allow all active reorganizations to complete and then terminate.

Environmental errors are categorized as errors related to the current availability of resources required to reorganize a single data set. Insufficient space during allocation of backup data sets is an example of an environmental error. FDRREORG will allow for up to 99 environmental errors. You can use the MAXENVERR keyword on the REORG statement to change this value.

System errors are categorized as any system abend or as any error that prevents FDRREORG from completing a reorganization. I/O errors, GETMAIN failures, and other system service failures are examples of system errors. FDRREORG will allow for up 99 system errors. You can use the MAXSYSERR keyword on the REORG statement to change this value.

**CANCEL
PROTECTION**

FDRREORG should never be canceled. Cancellation could result in data corruption or loss.

To protect against this, FDRREORG has automatic CANCEL protection. If the system operator or a user tries to cancel a FDRREORG step, FDRREORG will intercept the CANCEL and issue a FDRW99 message to the system operator asking for instructions. The operator can reply **I** to ignore the CANCEL and continue processing

S to stop after completing the data sets currently being reorganized

C to accept the CANCEL and terminate immediately. **This may result in data corruption or loss.**

The operator should be instructed to always respond S to the FDRW99 message, and FDRREORG will stop after processing the data set it is currently working on.

30.01 CONTINUED . . .

**IFANY and
IFALL
KEYWORDS**

In most cases, you would want a data set excluded from or selected for reorganization if any of the appropriate thresholds are met or exceeded. For those situations where a very stringent exclusion or selection process is desired, you may want all of the appropriate thresholds to be met or exceeded. The IFANY and IFALL keywords can be used to modify FDRREORG's evaluation of the following thresholds in just such a manner:

IAM-Enhanced	VSAM	PDS
OFULL	CASPLITR	PDSFULL
PCTTRECO	CISPLITR	PDSEXTENTS
PEFULL	PEUDATAR	

By default, FDRREORG will evaluate these thresholds as if the IFANY keyword was specified on the EXCLUDE or SELECT statement. If you want all of these thresholds to be met or exceeded before a data set is excluded or selected, use the IFALL keyword on the EXCLUDE or SELECT statement. If you specify any other criteria on your exclude or select statement, they will all have to be satisfied as well.

**DSN= and
CATDSN=
DIFFERENCES**

Although the end result is usually the same, your choice in using the DSN= versus the CATDSN= keyword can significantly impact FDRREORG's processing. Choosing between these two keywords is actually very simple once you understand the differences in FDRREORG's behavior when you choose one over the other.

The most significant difference between DSN and CATDSN relates to partitioned data sets. If you use CATDSN, uncataloged PDS's will never be selected. This does not apply to IAM or VSAM data sets because they must always be cataloged.

When CATDSN is used, FDRREORG scans the appropriate catalog(s) when the control cards are parsed and limits its volume processing to the volumes that match the filter or names you specify. Use of the VOL= keyword is optional and if specified is used to restrict the data sets to the volumes or volume groups you provide. The remaining selection criteria will not be evaluated until the volume is processed during the actual simulate or reorganization phase. Care should be taken when using the data set filtering capabilities because the data set names are kept in storage throughout simulate or reorg processing. Avoid coding statements such as CATDSN=** without a restrictive volume list. Doing so would cause a very large data set list to be built which would most probably result in an abnormal termination of FDRREORG. CATDSN=** also causes FDRREORG to read every available catalog which could take a considerable amount of time if there are a large number of user catalogs.

When DSN is used, a catalog scan does not occur. FDRREORG performs most of the exclusion and selection process when reading the volumes VTOC for IAM and PDS's, and the VVDS for VSAM. Catalog locates are issued for IAM and VSAM data sets for additional information after the data set and volume filtering is done. Because the primary source of information is the VTOC or VVDS, the VOL= keyword must be provided. Every volume that matches a volume or volume group in the volume list will be processed.

As a general rule, you should use the CATDSN keyword when the number of data sets that would be returned from the catalog scan are small, or the volume naming conventions in use at your installation do not translate well into volume groups. Use the DSN keyword when many data sets could be returned from the catalog scan, or you are more interested in processing particular volumes as opposed to specific data sets. If you have uncataloged partitioned data sets that you want to reorganize, you must use the DSN keyword.

SAMPLE FDRREORG OUTPUT OF PO DATA SETS

				CISR		CASR						RECORDS					
				%OFU	%PEU	TRKS	TRKS	PCT	NUM			TRKS	TRKS	PCT	NUM		
				DIRU	DIRA	USED	ALOC	USE	EXT	DIRU	DIRA	USED	ALOC	USE	EXT		
PAYROLL.TEST.LINKLIB	TS0001	P0	COMPRESSED	102	115	581	600	97	2	102	115	475	600	79	2		
SMPE.SMPPTS	DLIB01	P0	COMPRESSED	251	500	6334	6696	95	1	251	500	5725	6696	85	1		
USER.LOADLIB	TS0010	P0	COMPRESSED	436	1108	1061	1275	83	1	436	1108	932	1275	73	1		
USER2.JCLLIB	TS0015	P0	ALREADY COMPRESS	35	40	11	15	73	1	35	40	11	15	73	1		

↑

**DIRECTORY BLOCKS
USED AND ALLOCATED**

↑

**TRACKS USED
BEFORE FDRREORG**

↑

**TRACKS USED
AFTER FDRREORG**

Sample Report of reorganizing VSAM data sets. The backup will be created as a GDG. Only data sets with either CI or CA split ratios greater than 4 will be reorganized.

REORG	BACKUP=GDG														
SELECT	CATDSN=(PROD.** ,SMPE.**),DSTYPE=VSAM,ALWAYSBACKUP, CASPLITR>4,CISPLITR>4														
				CISR CASR						RECORDS					
				%OFU	%PEU	TRKS	TRKS	PCT	NUM	LOADED		TRKS	TRKS	PCT	NUM
DATASET NAME				DIRU	DIRA	USED	ALOC	USE	EXT	DIRU	DIRA	USED	ALOC	USE	EXT

PROD.ACCTPAY.MASTER	PROD01	VS	REORGANIZED	13	27	55	55	100	10	5125		45	55	82	10
PROD.ACCTPAY.ALTMSR	PROD01	VS	REORGANIZED	26	53	85	85	100	16	7367		50	85	59	16
PROD.HRIS.HISTORY	PROD02	VS	BACKED UP	0	0	450	480	100	1						
PROD.HRIS.EMPLOYEE.MASTER	PROD02	VS	REORGANIZED	49	47	95	95	100	18	10137		70	95	74	18
PROD.PAYROLL.MASTER	PROD03	VS	REORGANIZED	13	27	55	55	100	10	5226		45	55	82	10
PROD.PAYROLL.SUSPENSE	PROD03	VS	REORGANIZED	23	39	65	65	100	12	6124		45	65	69	12
SMPE.SP223.CSI	DLIB01	VS	REORGANIZED	5	16	1305	1350	97	1	99200		1125	1350	83	1
SMPE.SP422.CSI	DLIB01	VS	REORGANIZED	2	6	780	900	87	1	65222		735	900	82	1
SMPE.GLOBAL.CSI	DLIB01	VS	REORGANIZED	56	70	300	330	91	1	21805		150	330	45	1

The ALWAYSBACKUP keyword causes this data set to be chosen by selection criteria but did not qualify for reorganization.

CI AND CA SPLIT RATIO

NUMBER OF RECORDS LOADED

- 30-9 -

SAMPLE FDRREORG OUTPUT OF IAM DATA SETS

```
* REORG
* SELECT CATDSH=PROD.BILLING.** ,DSTYPE=IAM
```

DATASET NAME	VOLUME	AM	STATUS	CISR CASR		TRKS USED	TRKS ALLOC	PCT USE	NUM EXT	RECORDS LOADED		TRKS USED	TRKS ALLOC	PCT USE	NUM EXT
				DIRU	DIRA					DIRU	DIRA				
PROD.BILLING.HISTORY	PROD04	IA	REORGANIZED	93	0	542	570	95	1	5356	581	615	95	2	
PROD.BILLING.MASTER	PROD05	IA	REORGANIZED	100	0	7415	7500	99	8	78139	7309	7500	97	8	
PROD.BILLING.PAYED	PROD03	IA	REORGANIZED	81	10	300	300	100	1	3487	300	300	100	1	
PROD.BILLING.XCOLLECT	PROD05	IA	REORGANIZED	10	100	150	150	100	1	1853	140	150	93	1	

PERCENTAGE OF OVERFLOW USED PERCENTAGE OF PRIME EXTENSION USED NUMBER OF RECORDS LOADED

30.02 FDRREORG JOB CONTROL REQUIREMENTS AND DEFAULTS

To execute FDRREORG use the following JCL:

EXEC STATEMENT Specifies the program name (PGM=FDRREORG), region requirement (REGION), and optional PARM= operand.

Region requirements vary by the type of data sets being processed and the number of subtasks invoked. Generally, simulation requires 3 megabytes of storage. REORG usually requires 4 to 9 megabytes of virtual storage below the line. FDRREORG utilizes storage above the line whenever possible.

If a PARM field is specified, FDRREORG will use data specified as the first control statement, which must be a REORG, RECOVER or SIMREORG statement. If FDRCOPY is invoked from another program, you can pass control statements using IBM's convention for passing data from the PARM field; contact Innovation for assistance if you wish to do this.

STEPLIB or JOBLIB DD STATEMENT If required, must specify the load module library in which FDRREORG resides. You may also need to specify the name of the IAM load module library if you have IAM installed and it is not in LNKLIST. These must be APF authorized.

SYSPRINT DD STATEMENT Specifies the output message data set. It must be present and is usually a SYSOUT data set. Messages from FDRREORG's main task are written to this DD.

REORGPRD DD STATEMENT Specifies the output message data set. It must be present and is usually a SYSOUT data set. Messages from FDRREORG's subtasks are written to this DD.

REORGPRD DD STATEMENT Specifies the report data set. It must be present and is usually a SYSOUT data set.

SYSIN DD STATEMENT Specifies the control statement data set required for all functions. Usually an input stream or DD * data set.

30.03 FDRREORG REORG AND SIMULATE STATEMENTS

REORG REO	SIMULATE SIM	
,ALIASCHECK=YES NO	,DSNRETRY= <u>NO</u> RETRY ENQ WAIT	,MODE= <u>ALL</u> PARALLEL SINGLE
,ALWAYSBACKUP	,DSTYPE=(type1,..typen)	,MOVEAIX=YES <u>NO</u>
,BACKUP=TEMP PERM GDG	,EMPTYBACKUPS= <u>DELETE</u> KEEP	,MSGLEVEL=I W E
,BACKUPALLOC=SMS UNIT	,ENQERR=YES WRN NO	,MSGTIMESTAMP=YES NO
,BACKUPDATACLASS=dataclas	,IAMCOMPPERM=YES NO	,NODEFAULTS
,BACKUPEXPDT=yydddyyyyddd	,IAMCOMPTEMP=YES NO	,NONSMS=KEEP SMS
,BACKUPGROUP=groupname	,IAMDEFINE=YES NO	,NOREORG
,BACKUPINDEX=index	,IAMDISP= <u>OLD</u> SHR	,NOUPDATES=YES NO
,BACKUPMGMTCLASS=mgmtclas	,LASTAPEPREFIX=prefix	,OWNERSTRING=ownr
,BACKUPRETPD=nnn	,LASTAPE	,PDSDISP=SHR <u>OLD</u>
,BACKUPSTORCLASS=storclas	,LISTNOREORG=YES NO	,RUNTIME=nnnn
,BACKUPSTRING=(old,new)	,LOGALLOC=SMS UNIT	,SELTERR=YES WRN NO
,BACKUPUNIT=unitname	,LOGDATACLASS=dataclas	,SIMPDSCOMP
,BACKUPUNITS=n	,LOGMGMTCLASS=mgmtclas	,SMS=KEEP SMS
,BUILDEMPYAIX=YES <u>NO</u>	,LOGPREFIX=prefix	,SMSADDDVOL
,CKPTALLOC=SMS UNIT	,LOGSTORCLASS=storclas	,SORTRPT=YES NO
,CKPTDATACLASS=dataclas	,LOGUNIT=unitname	,STOPTIME=hhmm
,CKPTMGMTCLASS=mgmtclas	,MAXENQ=nnn	,SUBTASKABEND=CONT TERM
,CKPTPREFIX=prefix	,MAXENVERR=nnnnn	,UPDATEDPDS=YES NO
,CKPTSTORCLASS=storclas	,MAXFILE=nnn	,VSAMDEFINE=ALWAYS IFREQ NO
,CKPTUNIT=unitname	,MAXPARALLELBACKUPS=n	,VSAMDISP= <u>OLD</u> SHR
,CONVERTINDEX	MAXSYSERR=nnnnn	
,DATA=ALLOC <u>USED</u>	,MAXTASKS=nn	

30.03 CONTINUED . . .

REORG STATEMENT	Use this statement to specify that REORG is to reorganize selected data sets.										
SIMULATE STATEMENT	Use this statement to specify that REORG is to report on data sets that would have been processed if REORG had been specified.										
OPERANDS	The following operands may be specified with the REORG or SIMULATE statement.										
ALIASCHECK=ALIASC	<p>Specifies if a check for an existing catalog alias for the backup data set should be done prior to allocating a backup data set.</p> <p>YES – Check for an existing alias and fail the backup or reorganization if one does not exist.</p> <p>NO – Do not check for an existing alias. If an alias does not exist, the backup data set will be cataloged in the master catalog. This will fail if the job does not have authority to update the master catalog.</p> <p>The default comes from the FDRREORG options table which is shipped set to YES.</p>										
ALWAYSBACKUP ALWAYS B	<p>Specifies that a backup should always be taken of IAM or VSAM data sets that meet the general selection criteria but do not qualify for reorganization. When using this keyword, data sets will qualify for a backup when ALL CRITERIA EXCEPT THOSE LISTED BELOW are met:</p> <table> <tr> <td><u>IAM</u></td><td><u>VSAM</u></td></tr> <tr> <td>OFULL</td><td>CASPLITR</td></tr> <tr> <td>PCTTRECO</td><td>CISPLITR</td></tr> <tr> <td>PEFULL</td><td>FREESPACE</td></tr> <tr> <td>PEUDATAR</td><td></td></tr> </table> <p>This keyword can also be specified on the SELECT statement. Specify ALWAYSBACKUP at the SELECT level if you want to limit this option to specific data sets.</p>	<u>IAM</u>	<u>VSAM</u>	OFULL	CASPLITR	PCTTRECO	CISPLITR	PEFULL	FREESPACE	PEUDATAR	
<u>IAM</u>	<u>VSAM</u>										
OFULL	CASPLITR										
PCTTRECO	CISPLITR										
PEFULL	FREESPACE										
PEUDATAR											
BACKUP=	<p>Defines the default disposition of the backup data sets after the target data set has been successfully reorganized. The value specified on the SELECT statement that selects the data set takes precedence.</p> <p>TEMP – Backup data sets on disk will be deleted and uncataloged. Backup data sets on tape will be uncataloged. If the LASTAPE option is specified, BACKUP=PERM will be forced.</p> <p>PERM – Backup data sets on disk will not be deleted and uncataloged. Backup data sets on tape will not be uncataloged.</p> <p>GDG – Same as PERM except backup data sets will be a +1 generation data set. REORG will dynamically define any generation data groups that do not already exist. The high level qualifier of the dynamically generated generation data groups must be aliased to an ICF catalog. GDG's are not supported for parallel backups.</p> <p>The default is taken from the FDRREORG option table which is shipped set to TEMP.</p>										

30.03 CONTINUED . . .

**BACKUPALLOC=
BACKUPA**

Defines the allocation method to use when allocating backup data sets.

SMS – Allocate by SMS storage class.

UNIT – Allocate by unit name.

The default is taken from the FDRREORG option table which is shipped set to UNIT.

**BACKUPDATACLASS=
BACKUPD**

Defines the default SMS data class for backup data sets. This operand can not be specified with BACKUPUNIT. The value specified on the SELECT statement that selects the data set takes precedence. Your installation ACS routines may override this value.

The default for this keyword can be set in the FDRREORG option table.

**BACKUPEXPDT=
BACKUPE**

Defines the default expiration date to be used for all backup data sets. The value specified on the SELECT statement that selects the data set takes precedence. Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

**BACKUPGROUP=
BACKUPG**

Defines the default group name to be used to generate the names of the backup data sets. The number of characters specified will replace, left to right, characters in the target data set name. The value specified on the SELECT statement that selects the data set takes precedence. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used. See the description of the SELECT statement for examples.

You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

If MODE=PARALLEL, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

EXAMPLE: BACKUPGROUP=BKUP?

The generated backup data set name will be the name of the selected data set with BKUPA overlaying the first 5 characters of the first backup name, BKUPB will overlay the first 5 characters of the second backup name, and so on.

30.03 CONTINUED . . .

**BACKUPINDEX=
BACKUPI**

Defines the default pattern to be used to add or delete index levels when generating the backup data set name. REORG will use each index level specified in BACKUPINDEX in place of the original index level. If a period is specified without any characters following, the original index level will be copied to the backup data set name. If + is specified, the character following the + will be inserted into the backup data set name as a new index level. If ++ is specified, the characters following the ++ will be added to the end of the backup data set name as a new index level. If – is specified, the index level will be dropped from the backup data set name. The value specified on the SELECT statement that selects the data set takes precedence. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used. See the description of the SELECT statement for examples.

If MODE=PARALLEL, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

EXAMPLE: BACKUPINDEX=++BACKUP?

The generated backup data set name will be the name of the selected dataset with .BACKUPA appended to the name of the first backup, .BACKUPB will be appended to the name of second backup, and so on.

The default is taken from the FDRREORG option table which is shipped set to ++BACKUP.

**BACKUPMGMTCLASS=
BACKUPM**

Defines the default SMS management class for backup data sets. This operand can not be specified with BACKUPUNIT. The value specified on the SELECT statement that selects the data set takes precedence. Your installation ACS routines may override this value.

The default for this keyword can be set in the FDRREORG option table.

**BACKUPRETPD=
BACKUPR**

Defines the default retention period to be used for all backup data sets. This operand can not be specified with BACKUPEXPDT. The value specified on the SELECT statement that selects the data set takes precedence.

The default for this keyword can be set in the FDRREORG option table.

CONTINUED . . .

30.03 CONTINUED . . .

**BACKUPSTORCLASS=
BACKUPS**

Defines a valid SMS storage class which will be used to allocate SMS managed backup data sets when BACKUPALLOC is SMS. This operand can not be specified with BACKUPUNIT. The storage class specified with the REORG statement will be used as the default storage class. The storage class specified on the SELECT statement that selects the data set, if any, takes precedence. Your installation ACS routines may override this value.

The default for this keyword can be set in the FDRREORG option table.

**BACKUPSTRING=
BACKUPSTR**

Specifies an old string and a new string enclosed in parentheses to be used in generating the name of a backup data set. Only the first occurrence of the old string will be replaced by the new string. See the description of the SELECT statement for examples.

If MODE=PARALLEL, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

Example: BACKUPSTRING=(CLUSTER, BACKUP?)

The generated backup data set name will be the name of the selected data set with BACKUPA replacing CLUSTER in the name of the first backup, BACKUPB will replace CLUSTER in the name of the second backup, and so on.

BACKUPUNIT=

Defines a valid unit name containing tape or disk devices to be used for allocating backup data sets when BACKUPALLOC is UNIT. The unit name specified must not define a group of devices that contain a mixture of device classes (ie. 3480s and 3380s). It is allowed to use a unit name that defines a group of devices with a mixture of device models (ie. 3380s and 3390s). This operand can not be specified with BACKUPSTORCLAS.

The default is taken from the FDRREORG option table which is shipped set to SYSDA.

BACKUPUNITS=

For backup data sets on tape, this operand defines the number of units to be allocated for backup data sets for each task. For backup data sets on disk, this operand defines the minimum number of units to be allocated for each backup data set. REORG will dynamically increase the number of disk units for a disk backup data set to ensure that sufficient space is available to backup the target data set.

The default is taken from the FDRREORG option table which is shipped set to 1.

30.03 CONTINUED . . .

BUILDEPTYAIX=

Specifies if FDRREORG should perform a build index on the empty alternate indexes of base clusters selected for reorganization.

YES – Build empty alternate indexes.

NO – Do not build empty alternate indexes.

The default is NO.

**CKPTALLOC=
CKPTA**

Defines the allocation method to use when allocating the checkpoint data set.

SMS – Allocate by SMS storage class.

UNIT – Allocate by unit name.

The default is taken from the FDRREORG option table which is shipped set to UNIT.

**CKPTDATACLASS=
CKPTD**

Specifies the SMS data class to use for the checkpoint file when CKPTALLOC is SMS. This operand can not be specified with CKPTUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

**CKPTMGMTCLASS=
CKPTM**

Specifies the SMS management class to use for the checkpoint file when CKPTALLOC is SMS. This operand can not be specified with CKPTUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

**CKPTPREFIX=
CKPTP**

Specifies the high level qualifier to use when constructing the name of the checkpoint file. To make it possible for the RECOVER statement to find the checkpoint file for a failed REORG, the checkpoint data set is cataloged with a name of

&CKPTPREFIX.REORGCKP.jobname.Dccyyddd.Thhmss

You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

The default is taken from the FDRREORG option table which is shipped set to FDRREORG.

**CKPTSTORCLASS=
CKPTS**

Specifies the SMS storage class to use for the checkpoint file when CKPTALLOC is SMS. This operand can not be specified with CKPTUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

30.03 CONTINUED . . .

**CKPTUNIT=
CKPTU**

Specifies a unit name containing disk devices to use for the checkpoint file when CKPTALLOC is UNIT. This keyword can not be specified with CKPTDATACLASS, CKPTMGMT CLASS, or CKPTSTORCLASS.

The default is taken from the FDRREORG option table which is shipped set to SYSDA.

CONVERTINDEX

Directs FDRREORG to convert an imbedded index of a VSAM KSDS or AIX to noimbed if it is redefined with a dataclass that requires extended format. VSAM compaction or extended addressability require extended format. Without this keyword, FDRREORG will issue message FDRS63 and redefine the KSDS or AIX using the original define parameters.

DATA=

Specifies whether to use ALLOCATED or USED space for the CYLS and TRKS options of the EXCLUDE and SELECT statements.

ALLOC – Exclude/Select processing will use allocated space for CYLS or TRKS.

USED – Exclude/Select processing will use used space for CYLS or TRKS.

The default is USED.

**DSNRETRY=
DSNR**

Defines the default action to take for data sets that are not available when they are selected for processing. The value specified on the SELECT statement that selects the data set takes precedence.

NO – The data set is bypassed.

RETRY – The data set will be added to the task's retry queue and periodic attempts will be made to allocate the data set. If the data set does not become available before the task completes its processing, the data set is bypassed.

ENQ – The data set will be added to the task's retry queue and an ENQ for the data set will be left pending. If the ENQ is not satisfied before the task completes its processing, the data set is bypassed.

WAIT – The data set will be added to the task's retry queue and an ENQ for the data set will be left pending. The task will wait for all data sets queue'd with the wait option before terminating. An operator STOP (P) command can be used to shutdown FDRREORG normally if it is no longer desirable to wait for data sets

The default is NO. You should not override the default in a JES3 installation.

30.03 CONTINUED . . .

DSTYPE=

Specifies the type(s) of data sets to be considered for reorganization.

IAM – IAM data sets are eligible for reorganization.

VSAM – VSAM data sets are eligible for reorganization.

PDS – Partitioned data sets are eligible for reorganization (compression).

ALL – All data sets are eligible for reorganization.

PDS is ignored if MODE=PARALLEL.

The default is ALL.

EMPTYBACKUPS=

Specifies if FDRREORG should keep the backups of previously loaded IAM or VSAM data sets that are now empty. Specifying KEEP is useful if the backups created by FDRREORG are retained as application backups and the current generation of each backup should reflect the current contents of all the data sets within that application.

DELETE – Empty backup data sets are deleted.

KEEP – Empty backup data sets are kept.

The default is DELETE.

ENQERR=

Specifies if FDRREORG should set return code 8, 4 or 0 if a data set selected for reorganization is in use by another job or user.

YES – Set return code 8 if a data set is in use.

WRN – Set return code 4 if a data set is in use.

NO – Set return code 0 if a data set is in use.

**IAMCOMPPERM=
IAMCOMPP**

Specifies if FDRREORG should backup compressed IAM data sets in compressed format when BACKUP=GDG or BACKUP=PERM has been specified.

YES – Backups will be in compressed format.

NO – Backups will not be in compressed format.

The default is taken from the FDRREORG option table which is shipped set to NO.

NOTE: This feature requires IAM 6.3 or above and is only supported for files in enhanced format.

30.03 CONTINUED . . .

**IAMCOMPTEMP=
IAMCOMPT**

Specifies if FDRREORG should backup compressed IAM data sets in compressed format when BACKUP=TEMP has been specified.

YES – Backups will be in compressed format.

NO – Backups will not be in compressed format.

The default is taken from the FDRREORG option table which is shipped set to YES.

NOTE: This feature requires IAM 6.3 or above and is only supported for files in enhanced format.

**IAMDEFINE=
IAMDEF**

Specifies if FDRREORG should delete and define IAM data sets that will be reorganized.

YES – IAM data sets will be deleted and defined before reloading.

NO – IAM data sets will not be deleted and defined before reloading.

The default is taken from the FDRREORG option table which is shipped set to NO.

NOTE: This feature requires IAM 6.3 or above. Additionally, files defined with an earlier release of IAM will not be redefined by FDRREORG until they are first redefined with IAM 6.3 or above.

**IAMDISP=
IAMD**

Specifies the disposition REORG is to use when allocating IAM data sets to be reorganized.

OLD – IAM data sets will be allocated DISP=OLD.

SHR – IAM data sets will be allocated DISP=SHR.

The default and recommended value is OLD !

**LASTAPEPREFIX=
LASTAPEP**

Specifies the high level qualifier to be used when cataloging or locating the special LASTAPE catalog entries. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

The default is taken from the FDRREORG option table which is shipped set to FDRREORG.

LASTAPE

Requests that the REORG processor append new backup data sets to the tape volumes used in a prior run. At the end of a successful run, a data set will be cataloged for each task of the form

&LASTAPEPREFIX.jobname.TASKnn.LASTAPE .

The catalog entry identifies the last volume used and the next file sequence number to use. If the next file to be backed up would have gone to a new tape, no LASTAPE will be cataloged. Use of this option requires the BACKUP operand to be specified as either PERM or GDG. If BACKUP=TEMP is specified or defaulted, it will be forced to PERM.

30.03 CONTINUED . . .

**LISTNOREORG=
LISTNOR**

Defines the NOREORG list print option.

YES – The NOREORG list is printed for REORG or SIMULATE functions.

NO – The NOREORG list is not printed.

The default is taken from the FDRREORG option table which is shipped set to YES.

**LOGALLOC=
LOGA**

Defines the allocation method to use when allocating the log data set.

SMS – Allocate by SMS storage class.

UNIT – Allocate by unit name.

The default is taken from the FDRREORG option table which is shipped set to UNIT.

**LOGDATACLASS=
LOGD**

Specifies the SMS data class to use for the log file when LOGALLOC is SMS. This keyword can not be specified with LOGUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

**LOGMGMTCLASS=
LOGM**

Specifies the SMS management class to use for the log file when LOGALLOC is SMS. This keyword can not be specified with LOGUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

**LOGPREFIX=
LOGP**

Specifies the high level qualifier to use when constructing the name of the log file. To make it possible for the RECOVER statement to find the log file, the log data set is cataloged with a name of

&LOGPREFIX.REORGLOG.jobname.Dccyyddd.Thhmmss

You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

The default is taken from the FDRREORG option table which is shipped set to FDRREORG.

**LOGSTORCLASS=
LOGS**

Specifies the SMS storage class to use for the log file when LOGALLOC is SMS. This keyword can not be specified with CKPTUNIT. Your installation ACS routines may override this value.

The default can be set in the FDRREORG option table.

**LOGUNIT=
LOGU**

Specifies a unit name containing disk devices to use for the log file when LOGALLOC is UNIT. This keyword can not be specified with LOGDATACLASS, LOGMGMTCLASS, or LOGSTORCLASS.

The default is taken from the FDRREORG option table which is shipped set to SYSDA.

30.03 CONTINUED . . .

MAXENQ=	<p>Specifies the maximum number of outstanding ENQs allowed for the DSNRETRY ENQ and WAIT options for each volume processor task. When this limit is reached, data sets will be added to task retry queue as if the RETRY option was specified. The ENQ will be issued when the outstanding ENQ count falls below this maximum. Once the volume processor task has completed processing all volumes and is waiting for data sets queued with the WAIT option, no additional ENQs will be issued.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 999.</p>
MAXENVERR= MAXENV	<p>Specifies the maximum number of environmental errors allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set. No additional REORG or SIMULATE statements will be processed. Environmental errors are any backup or re-load failures not caused by a system abend. Insufficient space, or target data set not available, are examples of environmental errors.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 99.</p>
MAXFILE=	<p>Specifies the maximum number of backup files to place on a single tape volume.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 255.</p>
MAXPARALLELBACKUPS= MAXP	<p>Specifies the maximum number of concurrent backups for each data set when MODE=PARALLEL has been specified. For an individual data set, FDRREORG will additionally limit the number of concurrent data set backups to the number of USED volumes. You may specify a value between 2 and 9.</p> <p>The default is 2.</p>
MAXSYSERR= MAXS	<p>Specifies the maximum number of system abends allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set. No additional REORG or SIMULATE statements will be processed.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 99.</p>

30.03 CONTINUED . . .

**MAXTASKS=
MAXT**

Specifies the maximum number of concurrent volumes to be processed. You may specify any number from 1 to 15, inclusive. Please note that the actual number of subtasks possible is limited by the amount of available virtual storage. For **SIMULATE** statements with **MAXTASKS=15** a region size of 3 megabytes should be sufficient. For **REORG** statements with **MAXTASK=15**, a region size of 8.5 megabytes should be sufficient if only VSAM or IAM files are processed. If PDS data sets are compressed about 1.2 megabytes per **SUBTASK=** is required.

The default is taken from the FDRREORG option table which is shipped set to 4.

MODE=

Specifies the concurrent data set backup mode.

ALL – All data sets are processed.

PARALLEL – Only IAM and VSAM data sets with more than one USED volume are processed. FDRREORG will create multiple backup tasks to backup specific portions of the data set. Use the **MAXPARALLELBACKUPS** keyword to specify the number backups per data set.

NOTE – PDS's are bypassed in parallel mode.

SINGLE – Compresses PDS's and processes IAM and VSAM data sets not supported by parallel mode. This option is intended to be used to complement those jobs run with **MODE=PARALLEL**. All data sets which will not be processed in parallel mode will be processed in single mode.

The default is **ALL**.

MOVEAIX=

Specifies if alternate index data sets should be moved if the base cluster is moved.

YES – alternate index data sets are defined using the **NEWVOLSDATA** and **NEWVOLSINDEX** specified for the base cluster.

NO – alternate index data sets are defined on the original volumes.

The default is **NO**. **YES** will be forced if a **NONSMS** managed cluster is redefined as an **SMS** managed cluster or if an **SMS** managed cluster is redefined a **NONSMS** managed cluster.

**MSGLEVEL=
MSGL**

Specifies the lowest level message type to be displayed on the subtask print file which has the ddname **REORGPR**T.

I – Informational, warning, and error messages will be displayed.

W – Warning and error messages will be displayed

E – Only error messages will be displayed.

The default is taken from the FDRREORG option table which is shipped set to **I**.

30.03 CONTINUED . . .

**MSGTIMESTAMP=
MSGT**

Specifies whether messages written to REORGPRT will be suffixed with the current date, time, and the internal subtask id that issued the message. If YES is requested or defaulted, the LRECL of the REORGPRT file is increased to 151.

YES – Messages will be timestamped.

NO – Messages will not be timestamped.

The default is taken from the FDRREORG option table which is shipped set to YES.

NODEFAULTS

If specified, the default **selection criteria** from the FDRREORG option table will not be used when evaluating SELECT statements.

NOTE: You should specify this option if you want data sets to be selected based only on the keywords you specify.

NONSMS=

Specifies if FDRREORG's internal define should keep the current management status and volumes for NONSMS managed data sets or let the SMS ACS routines decide.

KEEP – NONSMS managed data sets are redefined as NONSMS managed data sets.

SMS – The installation SMS ACS routines will decide if the data set should be defined as a managed data set.

NOREORG

Specifies that selected data sets should be backed up but not reorganized.

This keyword can also be specified on the SELECT statement. Specify NOREORG at the SELECT level if you want to limit this option to specific data sets.

**NOUPDATES=
NOUP**

Specifies the action to take for IAM or VSAM data sets selected for reorganization that have had no adds, no deletes, and no updates.

YES – The data set will be reorganized.

NO – The data set will not be reorganized.

The default is taken from the FDRREORG option table which is shipped set to NO.

**OWNERSTRING=
OWN**

Specifies a string of up to four characters which will be used as an eye catcher in the ownerid of the data component of a VSAM KSDS. REORG will update the ownerid field of the data component's catalog entry to record information used to insure data set integrity and to identify data sets that were being processed by REORG when either a system failure occurred, or REORG was canceled or otherwise terminated. See the RECOVER statement for additional information.

The default is taken from the FDRREORG option table which is shipped set to FDR\$.

30.03 CONTINUED . . .

**PDSDISP=
PDSD**

Specifies the disposition REORG is to use when allocating partitioned data sets to be compressed.

OLD – Partitioned data sets will be allocated DISP=OLD.

SHR – Partitioned data sets will be allocated DISP=SHR.

The default value is OLD.

**RUNTIME=
RUN**

Defines the number of minutes that this REORG or SIMULATE statement is allowed to run. When this limit is reached, processing for all active data sets is completed and the next REORG or SIMULATE statement, if any, is processed.

SELTERR=

Specifies if FDRREORG should set a return code of 8 if no data sets are selected by REORG or SIMULATE. Due to the nature of the selection process and the selection criteria specified, this may be a natural occurrence.

YES – Set return code 8 if no data sets are selected.

WRN – Set return code 4 if no data sets are selected.

NO –Set return code 0 if no data sets are selected.

The default is taken from the FDRREORG option table which is shipped set to YES.

SIMPDSCOMP

If specified on a SIMULATE statement, requests that a simulated compression be performed on partitioned data sets. Please note that all selected partitioned data sets will be read to perform a simulated compression.

SMS=

Specifies if FDRREORG's internal define should keep the current management status and volumes for SMS managed data sets or let the SMS ACS routines decide.

KEEP – SMS managed data sets are redefined as SMS managed data sets on the same volumes. However, on ESA and OS/390, only the first volume will be the same.

SMS – The installation SMS ACS routines will decide if the data set should be defined as a managed data set. SMS will select the volumes used to define the data set.

NOTE: If the SMS ACS routines indicate that a currently SMS managed data set should not be redefined as SMS managed, FDRREORG will force SMS=KEEP if NEWVOLSDATA and NEWVOLINDEX were not specified.

30.03 CONTINUED . . .

SMSADDVOL

Directs FDRREORG to dynamically add an SMS dummy candidate volume for SMS managed data sets before starting the reload. If more space is required during the reload than is available on the current volume(s), the candidate volume will be converted to a real volume and the reload should be able to complete without an out of space error. If the volume is not used, FDRREORG will remove it.

NOTE – This feature is not supported for IAM data sets defined with guaranteed space.

**SORTRPT=
SORTR**

Specifies the sorting option for the report written to REORGRPT. Specifying YES requires that information related to each data set processed or selected be kept in storage until the function is completed.

YES – Report is produced at the end of a REORG, SIMULATE or RECOVER function in data set name and volser order.

NO – Report is produced as data sets are processed or selected.

The default is taken from the FDRREORG option table which is shipped set to NO.

**STOPTIME=
STOP**

Defines a time in hours and minutes (hhmm or hh.mm) that processing for this statement is to stop. When this time is reached, processing for all active data sets is completed and the next REORG or SIMULATE statement, if any, is processed.

**SUBTASKABEND=
SUBTASKA**

Specifies the action to take if a volume processor subtask abends.

CONT – Continue processing without the subtask.

TERM – Quiesce all active work and terminate.

The default is taken from the FDRREORG option table which is shipped set to CONT.

UPDATEDPDS=

Specifies the action to take for partitioned data sets that do not have a current backup. If the update indicator is on in a data sets format 1 DSCB, the data set is considered to not have a current backup.

YES – The data set will be compressed.

NO – The data set will not be compressed.

The default is taken from the FDRREORG option table which is shipped set to YES.

30.03 CONTINUED . . .

**VSAMDEFINE=
VSAMDEF**

Specifies when FDRREORG should delete and define VSAM KSDSs that will be reorganized.

ALWAYS – VSAM KSDSs are always deleted and defined before reloading.

IFREQ – Only VSAM KSDSs that cannot be reused will be deleted and defined before reloading.

NO – Disables delete and define of all VSAM KSDSs. KSDSs that cannot be reused will not be reorganized.

This keyword can also be specified on the SELECT statement.

The default is taken from the FDRREORG option table which is shipped set to IFREQ.

**VSAMDISP=
VSAMD**

Specifies the disposition REORG is to use when allocating VSAM data sets to be reorganized.

OLD – VSAM data sets will be allocated DISP=OLD.

SHR – VSAM data sets will be allocated DISP=SHR.

The default and recommended value is OLD !

30.04 FDRREORG EXCLUDE STATEMENT

EXCLUDE
EX

ALL DATA SETS

,CRDAYS=nnn	,IFANY
,CRDATE=yydddlyyyyddd	,LRDAYS=nnnn
,CYLS=nnnn	,LRDATE=yydddlyyyyddd
,DSN=(dsn1,...dsn)	,TRKS=nnnnn
,DSTYPE=(dstypl,...dstypn)	,VOL=(vol1,...voln)
,IFALL	

VSAM DATA SETS ONLY

,CASPLITR=nnn	,NCASPLITS=nnn
,CISPLITR=nnn	,NCISPLITS=nnn
,FREESPACE=nnn	

IAM DATA SETS ONLY

,OFULL=nn	,PEFULL=nn
,ORECS=nnnnn	,PEBLKS=nnnnn
,OVERFLOWINDEX=nnnnnnnnnn	,PEUDATAR=nnnnn
,PCTTRECO=nnn	

PDS's ONLY

,PDSEXTENTS=nn	,PDSFULL=nn
----------------	-------------

30.04 CONTINUED . . .

EXCLUDE STATEMENT	Use this statement to specify exclusion thresholds for REORG or SIMULATE processing. An EXCLUDE statement applies to all SELECT statements that follow it. If any of the thresholds specified for an EXCLUDE are satisfied, the data set or volume will be bypassed.
OPERANDS	The following operands may be specified with EXCLUDE to restrict REORG or SIMULATE processing.
CASPLITR=CA	<p>Defines the ratio of Control Area splits to every 100 Control Area's in a VSAM KSDS.</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>
CISPLITR=CI	<p>Defines the ratio of Control Interval splits to every 100 Control Interval's in a VSAM KSDS.</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>
CRDAYS=	<p>Defines the number of days since a file was created.</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>
CRDATE=	<p>Defines the day a file was created. Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>
CYLS=	<p>Defines the size of a data set in Cylinders. Allocated or used space will be used as requested by the DATA= operand of the REORG or SIMULATE statement.</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>
DSN=	Specifies a list of up to 50 data set names or data set filters to exclude. Refer to Section 80 for information on specifying a data set filter.
DSTYPE=DST	<p>Specifies the type(s) of data sets to be excluded.</p> <p>IAM – IAM data sets are excluded.</p> <p>VSAM – VSAM data sets are excluded.</p> <p>PDS – Partitioned data sets are excluded.</p> <p>ALL – All data sets are excluded.</p> <p>The default is ALL.</p>
FREESPACE=FREE SP	<p>Defines the VSAM CA or CI free space percent. This keyword is used to prevent reorganizations of VSAM KSDSs or AIXs that have very high free space percentages. To reduce splits for some files in online systems, some VSAM files are loaded with little or no free space and then altered to have a very high CA and/or CI freespace percent. If these files are processed by FDRREORG, the current freespace percentage is used resulting in a dramatic increase in space.</p> <p>This operand supports the following logical operators. =, ≠, >, >=, <, <=</p>

30.04 CONTINUED . . .

IFALL	<p>Specifies that all of the criteria that measure the level of disorganization within a data set need to be satisfied in order to exclude a data set. This keyword applies to the following keywords only:</p> <table><tr><td><u>IAM</u></td><td><u>VSAM</u></td><td><u>PDS</u></td></tr><tr><td>OFULL</td><td>CASPLITR</td><td>PDSFULL</td></tr><tr><td>PCTTRECO</td><td>CISPLITR</td><td>PDSEXTENTS</td></tr><tr><td>PEFULL</td><td></td><td></td></tr><tr><td>PEUDATAR</td><td></td><td></td></tr></table> <p>IFANY is the default.</p>	<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>	OFULL	CASPLITR	PDSFULL	PCTTRECO	CISPLITR	PDSEXTENTS	PEFULL			PEUDATAR		
<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>														
OFULL	CASPLITR	PDSFULL														
PCTTRECO	CISPLITR	PDSEXTENTS														
PEFULL																
PEUDATAR																
IFANY	<p>Specifies that any of the criteria that measure the level of disorganization within a data set need to be satisfied in order to exclude a data set. This keyword applies to the following keywords only:</p> <table><tr><td><u>IAM</u></td><td><u>VSAM</u></td><td><u>PDS</u></td></tr><tr><td>OFULL</td><td>CASPLITR</td><td>PDSFULL</td></tr><tr><td>PCTTRECO</td><td>CISPLITR</td><td>PDSEXTENTS</td></tr><tr><td>PEFULL</td><td></td><td></td></tr><tr><td>PEUDATAR</td><td></td><td></td></tr></table>	<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>	OFULL	CASPLITR	PDSFULL	PCTTRECO	CISPLITR	PDSEXTENTS	PEFULL			PEUDATAR		
<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>														
OFULL	CASPLITR	PDSFULL														
PCTTRECO	CISPLITR	PDSEXTENTS														
PEFULL																
PEUDATAR																
LRDAYS=	<p>Defines the number of days since a data set was last referenced.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>															
LRDATE=	<p>Defines the date that the data set was last referenced. Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>															
NCASPLITS=	Defines the number of Control Area splits.															
NCA	This operand supports the following logical operators. =,≠,>,>=,<,<=															
NCISPLITS=	Defines the number of Control Interval splits.															
NCI	This operand supports the following logical operators. =,≠,>,>=,<,<=															
OFULL=	Defines the percent of overflow area used in an IAM file in compatible format.															
OF	This operand supports the following logical operators. =,≠,>,>=,<,<=															
ORECS=	Defines the number of records that have been allocated for the IAM overflow area for files in compatible format.															
OR	This operand supports the following logical operators. =,≠,>,>=,<,<=															
OVERFLOWINDEX=	<p>Specifies the amount of memory in bytes required for the in storage IAM overflow index.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>															
PCTTRECO=	<p>Defines the percent of total records in an IAM file that are in the IAM overflow area.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>															
PDSFULL=	Defines the percent of allocated space used by a PDS.															
PDSF	This operand supports the following logical operators. =,≠,>,>=,<,<=															

30.04 CONTINUED . . .

PDSEXTENTS=	Defines the number of extents for a PDS.
PDSEX	This operand supports the following logical operators. =, ≠, >, >=, <, <=
PEBLKS=	Defines the number of blocks that have been allocated for the IAM prime extension for files in compatible format.
PEB	This operand supports the following logical operators. =, ≠, >, >=, <, <=
PEFULL=	Defines the percent of the prime extension area used in a IAM file in compatible format.
PEF	This operand supports the following logical operators. =, ≠, >, >=, <, <=
PEUDATAR=	Defines the ratio times 100 of used prime extension blocks to prime data blocks in an IAM file.
PEUD	This operand supports the following logical operators. =, ≠, >, >=, <, <=
TRKS=	Defines the size of data set in tracks. Allocated or used space will be used as specified by the DATA= operand of REORG or SIMULATE.
	This operand supports the following logical operators. =, ≠, >, >=, <, <=
VOL=	Specifies a list of up to 50 volumes or volume groups. A volume group is specified by coding an asterisk at the end of the volser prefix (ie. VOL=PROD*). Use VOL=* to specify all volumes.

```

SELECT      DSN=(dsn1,...dsnn)           ,VOL=(vol1,...voln)
S           CATDSN=(dsn1,...dsnn)       ,STORGRP=storgrp *
          ALLDSN
          * VOL= or STORGRP= required with ALLDSN or DSN

```

- 30-32 -

30.05 CONTINUED . . .

SELECT STATEMENT This statement defines the selection criteria to be used for REORG or SIMULATE processing. It is possible to override operands specified on the REORG statement for the backup data sets created for the data sets selected by a particular SELECT statement.

Many of the operands on the SELECT statement are used to select data sets to be reorganized. The operands which apply to each data set type are shown in the appropriately labeled box on the previous page.

Some of the operands, especially those in the box labeled "Backup Data Sets Only", control how the reorganization of a data set selected by this SELECT statement will be processed.

OPERANDS The following operands may be specified with SELECT to control REORG or SIMULATE processing.

ALLDSN Specifies that all data sets are to be processed regardless of data set name. This option is required if DSN= or CATDSN= is not specified and the intent is truly to process all data sets. You can specify CATDSN=** or DSN=** instead of using ALLDSN.

**ALWAYSBACKUP
ALWAYSB** Specifies that a backup should always be taken of data sets that meet the general selection criteria but do not qualify for reorganization.

When using this keyword, data sets will qualify for a backup when ALL CRITERIA EXCEPT THOSE LISTED BELOW are met:

IAM

OFULL
PCTTRECO
PEFULL
PEUDATAR

VSAM

CASPLITR
CISPLITR
FREESPACE

If ALWAYSBACKUP is specified on the REORG statement, it will automatically be in effect for all SELECT statements.

BACKUP= Overrides the value specified or defaulted on the REORG statement for the disposition of backup data sets after a successful reorganization.

TEMP – Backup data sets on disk will be deleted and uncataloged. Backup data sets on tape will be uncataloged. If the LASTAPE option was specified on the REORG statement, BACKUP=PERM will be forced.

PERM – Backup data sets on disk will not be deleted and uncataloged. Backup data sets on tape will not be uncataloged.

GDG – Same as PERM except backup data sets will be a +1 generation data set. REORG will dynamically define any generation data groups that do not already exist. The high level qualifier of the dynamically generated generation data groups must be aliased to an ICF catalog.

GDG's are not supported for parallel backups.

30.05 CONTINUED . . .

BACKUPDATACLASS=
BACKUPD Defines a valid SMS data class which will be used instead of the data class specified on the REORG statement to allocate SMS managed backup data sets. This operand will be ignored if BACKUPUNIT was specified on the REORG statement. Your installation ACS routines may override this value.

BACKUPEXPDT=
BACKUPE Defines the expiration date to be used for the backup data sets created for data sets selected by this SELECT statement. This operand can not be specified with BACKUPRETPD, Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

BACKUPGROUP=
BACKUPG Defines the group name to be used to generate the names of the backup data sets. The number of characters specified will replace, left to right, characters in the target data set name. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

EXAMPLE:

```
SELECT DSN=(CICS.**),BACKUPGROUP=BKUP
```

All backup data sets will have a data set name starting with BKUP.

```
SELECT DSN=(CICS.**),BACKUPGROUP=&RACFUID
```

If the user id associated with this job is BKUP, all backup data sets will have a data set name starting with BKUP.

If MODE=PARALLEL, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

EXAMPLE:

```
BACKUPGROUP=BKUP?
```

The generated backup data set name will be the name of the selected data set with BKUPA overlaying the first 5 characters of the first backup name, BKUPB will overlay the first 5 characters of the second backup name, and so on.

30.05 CONTINUED . . .

**BACKUPINDEX=
BACKUPI**

Defines the pattern to be used to add or delete index levels when generating the backup data set name. REORG will use each index level specified in BACKUPINDEX in place of the original index level. If a period is specified without any characters following, the original index level will be copied to the backup data set name. If + is specified, the character following the + will be inserted into the backup data set name as a new index level. If ++ is specified, the characters following the ++ will be added to the end of the backup data set name as a new index level. If – is specified, the index level will be dropped from the backup data set name. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

EXAMPLES:

```
SELECT DSN=(CICS.**),BACKUPINDEX=++BACKUP
```

All backup data sets names will be generated using the target data set name with an additional index level of .BACKUP added to the end of the data set name.

```
SELECT DSN=(CICS.**),BACKUPINDEX=CICSBKUP.
```

All backup data sets names will be generated using the target data set name with the high level index of CICS replaced with CICSBKUP.

```
SELECT DSN=(CICS.**),BACKUPINDEX=CICSBKUP.++BACKUP
```

All backup data sets names will be generated using the target data set name with the high level index of CICS replaced with CICSBKUP and .BACKUP added as an additional index level at the end of the data set name.

The default is ++BACKUP or the value specified on the REORG statement.

```
SELECT DSN=(CICS.**),BACKUPINDEX=&RACFUID.++BACKUP
```

If the user id associated with this job is CICSBKUP, all backup data sets names will be generated using the target data set name with the high level index of CICS replaced with CICSBKUP and .BACKUP added as an additional index level at the end of the data set name.

The default is ++BACKUP or the value specified on the REORG statement.

If MODE=PARALLEL, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

EXAMPLE:

```
BACKUP INDEX=++BACKUP?
```

The generated backup dataset name will be the name of the selected dataset with .BACKUPA appended to the name of the first backup, .BACKUPB will be appended to the name of second backup, and so on.

30.05 CONTINUED . . .

BACKUPMGMTCLASS= Defines a valid SMS management class which will be used instead of the management class specified on the REORG statement to allocate SMS managed backup data sets. This operand will be ignored if **BACKUPM** **BACKUPUNIT** was specified on the REORG statement. Your installation ACS routines may override this value.

BACKUPRETPD= Defines the retention period to be used for the backup data sets **BACKUPR** created for data sets selected by this SELECT statement. This operand can not be specified with **BACKUPEXPDT**.

BACKUPSTORCLASS= Defines a valid SMS storage class which will be used instead of the **BACKUPS** storage class specified on REORG statement to allocate SMS managed backup data sets. This operand will be ignored if **BACKUPUNIT** was specified on the REORG statement. Your installation ACS routines may override this value.

BACKUPSTRING= Specifies an old string and a new string enclosed in parentheses to be used in generating the name of a backup data set. Only the first occurrence of the old string will be replaced by the new string. If the old string is not found, a valid backup data set name can not be generated and the selected data set will be bypassed. To ensure that this does not occur, always include the old string in the names specified via the DSN or CATDSN keywords.

EXAMPLE:

```
SELECT DSN=(CICS.VSAMCLUS.**),
      BACKUPSTRING=(VSAMCLUS,BKUPCLUS)
```

The string 'VSAMCLUS' will be replaced by the string 'BKUPCLUS' in the name of all backup data sets.

If **MODE=PARALLEL**, you must place a single question mark somewhere in the specified string. The question mark is replaced by a letter starting with A to identify the multiple backups taken in parallel mode.

EXAMPLE:

```
BACKUPSTRING=(CLUSTER, BACKUP?)
```

The generated backup dataset name will be the name of the selected dataset with **BACKUPA** replacing **CLUSTER** in the name of the first backup, **BACKKUPB** will replace **CLUSTER** in the name of second backup, and so on.

CASPLITR=
CA

Defines the ratio of Control Area splits to every 100 Control Area's in a VSAM KSDS.

This operand supports the following logical operators. =, <, >, >=, <=, <=

The default is >= the value in the FDRREORG option table which is shipped set to 10.

30.05 CONTINUED . . .

CATDSN=	Specifies a list of up to 50 data set names or 50 data set filters. This operand can not be specified if DSN is specified. When this operand is used, the system catalog(s) are searched for all data sets that match the specified data set names or filters and reside on the volumes provided in VOL= operand (if specified). The volume processor subtasks will only process the volumes for the selected data sets. To get a list of all data set names returned from the catalog search, use PCATDSN instead of CATDSN. Refer to section 80 for information on specifying a data set filter.
CISPLITR= CI	<p>Defines the ratio of Control Interval splits to every 100 Control Interval's in a VSAM KSDS.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 10.</p>
CRDAYS=	<p>Defines the number of days since a file was created.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>
CRDATE=	<p>Defines the day a file was created. Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>
CYLS=	<p>Defines the size of a data set in Cylinders. Allocated or used space will be used as requested by the DATA= operand of the REORG or SIMULATE statement.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>
%DIRFREE=	Specifies the percentage of free PDS directory blocks required. The PDS directory will be expanded to this value if required.
DIRBLKS=	Specifies the number of PDS directory blocks to be added.
DSN=	Specifies a list of up to 50 data set names or 50 data set filters. This operand can not be specified if CATDSN is specified. Data set selection is not done until the volumes provided in the VOL= operand are processed. Refer to Section 80.14 for information on specifying a data set filter.

30.05 CONTINUED . . .

**DSTYPE=
DST**

Specifies the type(s) of data sets to be considered for reorganization.

IAM – IAM data sets are eligible for reorganization.

VSAM – VSAM data sets are eligible for reorganization.

PDS – Partitioned data sets are eligible for reorganization (compression).

ALL – All data sets are eligible for reorganization.

PDS is ignored if MODE=PARALLEL.

ALL is the default value unless specified on the REORG or SIMULATE statement.

**DSNRETRY=
DSNR**

Defines the action to take for data sets that are not available when they are selected for processing.

NO – The data set is bypassed.

RETRY – The data set will be added to the task's retry queue and periodic attempts will be made to allocate the data set. If the data set does not become available before the task completes its processing, the data set is bypassed.

ENQ – The data set will be added to the task's retry queue and an ENQ for the data set will be left pending. If the ENQ is not satisfied before the task completes its processing, the data set is bypassed.

WAIT – The data set will be added to the task's retry queue and an ENQ for the data set will be left pending. The task will wait for all data sets queue'd with the wait option before terminating. An operator STOP (P) command can be used to shutdown FDRREORG normally if it is no longer desirable to wait for data sets.

The default is NO or the value specified on the REORG statement. You should not override the default in a JES3 installation.

**FREESPACE=
FREESP**

Defines the VSAM CA or CI free space percent. This keyword is used to prevent reorganizations of VSAM KSDSs or AIXs that have very high free space percentages. To reduce splits for some files in online systems, some VSAM files are loaded with little or no free space and then altered to have a very high CA and/or CI freespace percent. If these files are processed by FDRREORG, the current freespace percentage is used resulting in a dramatic increase in space.

This operand supports the following logical operators. =,≠,>,>=,<,<=

The default is < the value in the FDRREORG option table which is shipped set to 50.

**GDGEMPTY
GDGEMP**

Specifies that dynamically defined GDG's will be defined with the EMPTY attribute. Without this option, all dynamically defined GDG's are defined with the NOEMPTY attribute.

**GDGLIMIT=
GDGLIM**

Defines the maximum number of generation data sets that can be cataloged for a dynamically defined GDG.

The default is 5.

30.05 CONTINUED . . .

**GDGNOSCRATCH
GDGNOSCR**

Specifies that dynamically defined GDG's will be defined with the NOSCRATCH attribute. Without this option, all dynamically defined GDG's are defined with the SCRATCH attribute.

**IAMDEFINE=
IAMDEF**

Specifies if FDRREORG should delete and define IAM data sets that will be reorganized.

YES – IAM data sets will be deleted and defined before reloading.

NO – IAM data sets will not be deleted and defined before reloading.

The default is the value specified on the REORG or SIMULATE card.

Note – This feature requires IAM 6.3 or above. Additionally, files defined with an earlier release of IAM will not be redefined by FDRREORG until they are first redefined with IAM 6.3 or above.

IAMGROW=

Specifies the percent to increase the size of single volume IAM data sets that are redefined.

IFALL

Specifies that all of the criteria that measure the level of disorganization within a data set need to be satisfied in order to select a data set. This keyword applies to the following keywords only:

<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>
OFULL	CASPLITR	PDSFULL
PCTTRECO	CISPLITR	PDSEXTENTS
PEFULL		
PEUDATAR		

IFANY is the default.

IFANY

Specifies that any of the criteria that measure the level of disorganization within a data set need to be satisfied in order to select a data set. This keyword applies to the following keywords only:

<u>IAM</u>	<u>VSAM</u>	<u>PDS</u>
OFULL	CASPLITR	PDSFULL
PCTTRECO	CISPLITR	PDSEXTENTS
PEFULL		
PEUDATAR		

IFANY is the default.

LRDAYS=

Defines the number of days since a data set was last referenced.

This operand supports the following logical operators. =,≠,>,>=,<,<=

LRDATE=

Defines the date that the data set was last referenced. Specify the date in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

This operand supports the following logical operators. =,≠,>,>=,<,<=

30.05 CONTINUED . . .

**MAXPRIBLKS=
MAXP=**

Specifies the maximum number of blocks that FDRREORG will use for the primary space allocation when allocating backup data sets on disk. If the estimated total blocks calculated by FDRREORG exceeds this value, FDRREORG will increase the unit count by 1 for each increment of this value.

If the volumes used for backup data sets are extremely fragmented, reducing this value should reduce the number of space related dynamic allocation failures. Keep in mind that additional units may be required to satisfy the allocation request which may not be available. If very large data sets are reorganized and there are a limited number of volumes in the pool used for backup data set allocations, increasing this value should reduce the number of unit related dynamic allocation failures. Keep in mind that requesting a larger primary allocation could cause allocations to fail because of insufficient space.

The default is 50000 blocks.

**NCASPLITS=
NCA**

Defines the number of Control Area splits.

This operand supports the following logical operators. =,≠,>,>=,<,<=

**NCISPLITS=
NCI**

Defines the number of Control Interval splits.

This operand supports the following logical operators. =,≠,>,>=,<,<=

NEWALL

This keyword is required to redefine multiple VSAM KSDSs with new space, volumes, or freespace via a single select statement that specifies multiple data set names or uses a data set mask. Without this keyword, FDRREORG will treat the combination of new space, volumes, or freespace as a keyword error to prevent accidentally redefining many VSAM KSDSs with the same space, volumes or freespace.

**NEWCYLSDATA=
NEWCYLSD**

Specifies a new primary and optional secondary space in cylinders to be used for the data component of a VSAM KSDS or an IAM file, if it is redefined.

EXAMPLES:

```
NEWCYLSDATA=( 100 , 10 )
```

The redefined data component or IAM file will have 100 cylinders of primary space and 10 cylinders of secondary space.

```
NEWCYLSDATA=100
```

The redefined data component or IAM file will have 100 cylinders of primary space and no secondary space.

30.05 CONTINUED . . .

**NEWCYLSINDEX=
NEWCYLSI**

Specifies a new primary and optional secondary space in cylinders to be used for the index component of a VSAM KSDS if it is redefined.

EXAMPLES:

```
NEWCYLSINDEX=( 2 , 1 )
```

The redefined index component will have 2 cylinders of primary space and 1 cylinder of secondary space.

```
NEWCYLSINDEX=2
```

The redefined index component will have 2 cylinders of primary space and no secondary space.

NEWDATACLASS=

Specifies the name of a valid SMS data class to be used to redefine a VSAM cluster and its alternate indexes if any, or an IAM file. The installation SMS ACS routine may override this value.

**NEWFREESPACE=
NEWFREESP**

Specifies a new CI and CA freespace percentage to be used for the data component of a VSAM KSDS or an IAM file, if it is redefined.

EXAMPLES:

```
NEWFREESPACE=( 10 , 20 )
```

The redefined data component or IAM file will have 10 percent CI freespace and 20 percent CA freespace.

NEWMGMTCLASS=

Specifies the name of a valid SMS management class to be used to redefine a VSAM cluster and its alternate indexes if any, or an IAM file. **The installation SMS ACS routine may override this value.**

**NEWRECSDATA=
NEWRECS**

Specifies a new primary and optional secondary space in records to be used for the data component of a VSAM KSDS or an IAM file, if it is redefined.

EXAMPLES:

```
NEWRECSDATA=( 10000 , 1000 )
```

The redefined data component or IAM file will have 10000 records of primary space and 1000 records of secondary space.

```
NEWRECSDATA=10000
```

The redefined data component or IAM file will have 10000 records of primary space and no secondary space.

NOTE: VSAM will convert the number of records specified into an appropriate number of cylinders or tracks.

30.05 CONTINUED . . .

**NEWRECSINDEX=
NEWRECSI**

Specifies a new primary and optional secondary space in records to be used for the index component of a VSAM KSDS if it is redefined.

EXAMPLES:

```
NEWRECSINDEX=( 500 , 50 )
```

The redefined index component will have 500 records of primary space and 50 records of secondary space.

```
NEWCYLSINDEX=500
```

The redefined index component will have 500 records of primary space and no secondary space.

NOTE: VSAM will convert the number of records specified into an appropriate number of cylinders or tracks.

NEWSTORCLASS=

Specifies the name of a valid SMS storage class to be used to redefine a VSAM cluster and its alternate indexes if any, or an IAM file. **The installation SMS ACS routine may override this value.**

**NEWTRKSDATA=
NEWTRKSD**

Specifies a new primary and optional secondary space in tracks to be used for the data component of a VSAM KSDS or an IAM file if it is redefined.

EXAMPLES:

```
NEWTRKSDATA=( 100 , 10 )
```

The redefined data component or IAM file will have 100 tracks of primary space and 10 tracks of secondary space.

```
NEWTRKSDATA=100
```

The redefined data component or IAM file will have 100 tracks of primary space and no secondary space.

**NEWTRKSINDEX=
NEWTRKSI**

Specifies a new primary and optional secondary space in tracks to be used for the index component of a VSAM KSDS if it is redefined.

EXAMPLES:

```
NEWTRKSINDEX=( 2 , 1 )
```

The redefined index component will have 2 tracks of primary space and 1 track of secondary space

```
NEWTRKSINDEX=2
```

The redefined index component will have 2 tracks of primary space and no secondary space.

**NEWVOLSDATA=
NEWVOLSD**

Specifies up to 20 volumes to be used for the data component of a VSAM KSDS or an IAM file, if it is redefined. If the data set is SMS managed, SMS=SMS will be forced for the define. SMS may or may not use the volumes specified.

EXAMPLES:

```
NEWVOLSDATA=( MYVOL 1 , MYVOL 2 )
```

30.05 CONTINUED . . .

**NEWVOLINDEX=
NEWVOLSI**

Specifies up to 10 volumes to be used for the index component of a VSAM KSDS if it is redefined. If the data set is SMS managed, SMS=SMS will be forced for the define. SMS may or may not use the volumes specified.

EXAMPLES:

```
NEWVOLINDEX=(MYVOL1,MYVOL2)
```

NOREVERT

Specifies that FDRREORG should not revert to the original define parameters if a VSAM KSDS redefine fails. Without this keyword, FDRREORG will always redefine a VSAM KSDS with the original space, volumes, and freespace parameters, if the define fails and any of these parameters were changed via the appropriate NEWxxxx keyword(s).

NOREORG

Specifies that selected data sets should be backed up but not reorganized.

If NOREORG is specified on the REORG statement, it will automatically be in effect for all SELECT statements.

**NOUPDATES=
NOUP**

Specifies the action to take for IAM or VSAM data sets selected for reorganization that have had no adds, no deletes and no updates.

YES – The data set will be reorganized.

NO – The data set will not be reorganized.

The default is the value specified on the REORG or SIMULATE statement.

**OFULL=
OF**

Defines the percent of overflow area used in an IAM file in compatible format.

This operand supports the following logical operators. =,≠,>,>=,<,<=

The default is >= the value in the FDRREORG option table which is shipped set to 80.

**ORECS=
OR**

Defines the number of records that have been allocated for the IAM overflow area for a file in compatible format.

This operand supports the following logical operators. =,≠,>,>=,<,<=

OVERFLOWINDEX=

Specifies the amount of memory in bytes required for the in storage IAM overflow index.

This operand supports the following logical operators. =,≠,>,>=,<,<=

The default is >= the value in the FDRREORG option table which is shipped set to 1048576 (1 megabyte).

**PCTOFKEEPMULTIPLE=
PCTOFKEEPM**

Specifies the percentage of IAM independent overflow to be retained for files in compatible format that were not loaded via a single record load.

The default is 100.

30.05 CONTINUED . . .

PCTOFKEEPSINGLE= PCTOFKEEPS	<p>Specifies the percentage of IAM independent overflow to be retained for files in compatible format that were loaded via a single record load.</p> <p>The default is 100.</p>
PCTPEKEEPMULTIPLE= PCTPEKEEPM	<p>Specifies the percentage of IAM prime extension blocks to be retained for files in compatible format that were not loaded via a single record load.</p> <p>The default is 100.</p>
PCTPEKEEPSINGLE= PCTPEKEEPS	<p>Specifies the percentage of IAM prime extension blocks to be retained for files in compatible format that were loaded via a single record load.</p> <p>The default is 100.</p>
PCTTRECO=	<p>Defines the percent of total records in an IAM file that are in the IAM overflow area.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 10.</p>
PDSFULL= PDSF	<p>Defines the percent of allocated space used by a PDS.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 90.</p>
PDSEXTENTS= PDSEX	<p>Defines the number of extents for a PDS. This operand supports the following logical operators. =,≠,>,>=,<,<= </p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 17.</p>
PEBLKS= PEB	<p>Defines the number of blocks that have been allocated for the IAM prime extension for a file in compatible format.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p>
PEFULL= PEF	<p>Defines the percent of the prime extension area used in a IAM file in compatible format.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 100.</p>
PEUDATAR= PEUD	<p>Defines the ratio times 100 of used prime extension blocks to prime data blocks in an IAM file.</p> <p>This operand supports the following logical operators. =,≠,>,>=,<,<=</p> <p>The default is >= the value in the FDRREORG option table which is shipped set to 100.</p>

30.05 CONTINUED . . .

PRISPACEPCT= PRISP	Defines the percentage to be used to calculate the primary space allocation for a backup data set that will be on disk. This percentage is applied to total used space for the target data set. The default is 60.
SECSPACEPCT= SECSP	Defines the percentage to be used to calculate the secondary space allocation for a backup data set that will be on disk. This percentage is applied to the primary space allocation used for the backup data set. The default is 25.
STORGRP=	STORGRP= can be used only on systems active with SMS; it will select all volumes in the specified SMS storage group.
TRKS=	Defines the size of a data set in tracks. Allocated or used space will be used as specified by the DATA= operand of REORG or SIMULATE. This operand supports the following logical operators. =,≠,>,>=,<,<=
VOL=	Specifies a list of up to 50 volumes or volume groups. A volume group is specified by coding an asterisk at the end of the volser prefix (i.e. VOL=PROD*). Use VOL=* to specify all volumes. NOTE: VOL= and STORGRP= are mutually exclusive.
VSAMDEFINE= VSAMDEF	Specifies when FDRREORG should delete and define VSAM KSDSs that will be reorganized. ALWAYS – VSAM KSDSs are always deleted and defined before reloading. IFREQ – Only VSAM KSDSs that cannot be reused will be deleted and defined before reloading. NO – Disables delete and define of all VSAM KSDSs. KSDSs that cannot be reused will not be reorganized. The default is the value specified on the REORG or SIMULATE card.
VSAMGROW=	Specifies the percent to increase the size of VSAM single volume KSDSs that are redefined.

30.06 FDRREORG RECOVER STATEMENT

RECOVER	BUILDEMPYAIX=YES <u>NO</u>	,NEWFREESPACE=nnn
REC	,CKPTPREFIX=prefix	,NEWMGMTCLASS=mgmtclas
	,CONVERTINDEX	,NEWRECSDATA=(nnnnn,nnnnn)
	,DSN=(dsn1,...dsnn)	,NEWRECSINDEX=(nnnnn,nnnnn)
	,ENQERR=YES NO	,NEWSTORCLASS=storclas
	,EXDSN=(dsn1,...dsnn)	,NEWTRKSDATA=(nnnnn,nnnnn)
	,IAMDISP= <u>OLD</u> SHR	,NEWTRKSINDEX=(nnnnn,nnnnn)
	,JOBNAME=jobname	,NEWVOLSDATA=(vol1,...voln)
	,LOGPREFIX=prefix	,NEWVOLSINDEX=(vol1,...voln)
	,MAXENVERR=nnn	,NOCKPT
	,MAXSYSERR=nnn	,NOLOG
	,MAXTASKS=nn	,NONSMS=KEEP SMS
	,MOVEAIX=YES <u>NO</u>	,SELTERR=YES NO
	,MSGLEVEL=I W E	,SMS=KEEP SMS
	,MSGTIMESTAMP=YES NO	,SMSADDVOL
	,NEWCYLSDATA=(nnnnn,nnnnn)	,SORTRPT=YES NO
	,NEWCYLSINDEX=(nnnnn,nnnnn)	,VSAMDISP= <u>OLD</u> SHR
	,NEWDATACLASS=dataclas	

RECOVER This statement is used to complete REORG processing for data sets that were not completely re-loaded from a prior run. If you use the JOBNAME keyword, FDRREORG will automatically allocate the appropriate checkpoint and log file data sets. Otherwise, the checkpoint and log files must be provided via DD statements. The checkpoint file is used to recover data sets that were being reorganized when FDRREORG abnormally terminated. Any data sets that were being re-loaded at this time are recovered. Data sets that had not completed the backup phase are bypassed. The log file is used to recover individual data sets that did not complete re-load processing. These data sets usually require some form of manual intervention before recovery will be successful. The most common cause will probably be insufficient space. This occurs because the data sets most likely selected for REORG processing may exhaust the allocated space if a significant number of records have been inserted into the files freespace area's. When the data set is reloaded, these records will be moved out of the files freespace areas which results in the increased space requirements.

NOTE – see [Section 30.12](#) for the FDRREORG Recover Examples.

30.06 CONTINUED . . .

OPERANDS The following operands may be specified with RECOVER to control recovery processing.

- BUILDEMPYAIX=** Specifies if FDRREORG should perform a build index on the empty alternate indexes of base clusters selected for reorganization. The default is NO.
- YES** – Build empty alternate indexes.
- NO** – Do not build empty alternate indexes.
- CKPTPREFIX=**
CKPTP Specifies the high level qualifier used when the checkpoint file was allocated. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used. The default is taken from the FDRREORG option table which is shipped set to FDRREORG.
- CONVERTINDEX** Directs FDRREORG to convert an imbedded index of a VSAM KSDS or AIX to NOIMBED if it is redefined with a dataclass that requires extended format. VSAM compaction or extended addressability require extended format. Without this keyword, FDRREORG will issue message FDRS63 and redefine the KSDS or AIX using the original define parameters.
- DSN=** Specifies a list of up to 50 data set names or 50 data set filters to recover. If this keyword is not provided, all data sets that qualify for recovery will be processed. This keyword can not be specified with EXDSN. Refer to [section 80](#) for information specifying a data set filter.
- The default is DSN=**
- ENQERR=** Specifies if FDRREORG should set return code 8 if a data set selected for recovery is in use by another job or user.
- YES** – Set return code 8 if a data set is in use.
- NO** – Do not set return code 8 if a data set is in use.
- EXDSN=** Specifies a list of up to 50 data set names or 50 data set filters to exclude from recovery processing. This keyword can not be specified with DSN. Refer to [Section 80](#) for information specifying a data set filter.
- IAMDISP=** Specifies the disposition REORG is to use when allocating IAM data sets to be recovered.
- OLD** – IAM data sets will be allocated DISP=OLD.
- SHR** – IAM data sets will be allocated DISP=SHR.
- The default and recommended value is OLD !

30.06 CONTINUED . . .

JOBNAME= JOB	<p>Specifies the name of the failed job to recover. FDRREORG will automatically allocate the checkpoint and log files from the specified job and recover all data sets that qualify for recovery. If this keyword is not specified, you MUST provide either a REORGCKP and/or REORGLOG DD statement to define the checkpoint and log files.</p> <p>Please note that depending on the type of failure that occurred, both of these files may not exist. The checkpoint file will only exist if FDRREORG was abnormally terminated, or if a volume processor subtask abended and was unable to log the name of the active data set to the log file. The log file will only exist if there were unsuccessful reloads. FDRREORG handles this automatically if you use the JOBNAME keyword. If you provide the DD statements via JCL, you must provide the appropriate NOCKPT or NOLOG keywords.</p>
LOGPREFIX= LOGP	<p>Specifies the high level qualifier used when the log file was allocated. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.</p> <p>The default is taken from the FDRREORG option table which is shipped set to FDRREORG.</p>
MAXENVERR= MAXENV	<p>Specifies the maximum number of environmental errors allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set.</p> <p>The default is 99.</p>
MAXSYSERR= MAXS	<p>Specifies the maximum number of system abends allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 99.</p>
MAXTASKS= MAXT	<p>Specifies the maximum number of concurrent data sets to be recovered. You may specify any number from 1 to 15, inclusive. If the backup files created by the failed REORG are on tape, there may be contention for the tape volumes if there were multiple failures from the same task.</p> <p>The default is taken from the FDRREORG option table which is shipped set to 4.</p>
MOVEAIX=	<p>Specifies if alternate index data sets should be moved if the base cluster is moved.</p> <p>YES – alternate index data sets are defined using the NEWVOLS specified for the base cluster.</p> <p>NO – alternate index data sets are defined on the original volumes.</p> <p>The default is NO. YES will be forced if a NONSMS managed cluster is redefined as an SMS managed cluster or if an SMS managed cluster is redefined a NONSMS managed cluster.</p>

30.06 CONTINUED . . .

MSGLEVEL=MSGL	<p>Specifies the lowest level message type to be displayed on the subtask print file which has the ddname REORGPRT.</p> <p>I – Informational, warning, and error messages will be displayed.</p> <p>W – Warning and error messages will be displayed</p> <p>E – Only error messages will be displayed.</p> <p>The default is taken from the FDRREORG option table which is shipped set to I.</p>
MSGTIMESTAMP=MSGT	<p>Specifies whether messages written to REORGPRT will be suffixed with the current date, time, and the internal subtask id that issued the message. If YES is requested or defaulted, the LRECL of the REORGPRT file is increased to 151.</p> <p>YES – Messages will be timestamped.</p> <p>NO – Messages will not be timestamped.</p> <p>The default is taken from the FDRREORG option table which is shipped set to YES.</p>
NEWCYLSDATA=NEWCYLSD	<p>Specifies a new primary and optional secondary space in cylinders to be used for the data component of a VSAM KSDS or an IAM file, if it is redefined.</p> <p>Examples:</p> <pre>NEWCYLSDATA=(100 , 10)</pre> <p>The redefined data component or IAM file will have 100 cylinders of primary space and 10 cylinders of secondary space.</p> <pre>NEWCYLSDATA=100</pre> <p>The redefined data component or IAM file will have 100 cylinders of primary space and no secondary space.</p>
NEWCYLSINDEX=NEWCYLSI	<p>Specifies a new primary and optional secondary space in cylinders to be used for the index component of a VSAM KSDS if it is redefined.</p> <p>Examples:</p> <pre>NEWCYLSINDEX=(2 , 1)</pre> <p>The redefined data component will have 2 cylinders of primary space and 1 cylinders of secondary space.</p> <pre>NEWCYLSINDEX=2</pre> <p>The redefined index component will have 2 cylinders of primary space and no secondary space.</p>
NEWDATACLASS=	<p>Specifies the name of a valid SMS data class to be used to redefine a VSAM cluster and its alternate indexes if any or an IAM file. Your installation ACS routines may override this value.</p>

30.06 CONTINUED . . .

NEWFREESPACE= Specifies a new CI and CA freespace percentage to be used for the data
NEWFREESP component of a VSAM KSDS or an IAM file, if it is redefined.

Examples:

```
NEWFREESPACE=( 10 , 20 )
```

The redefined data component or IAM file will have 10 percent CI freespace and 20 percent CA freespace.

NEWMGMTCLASS= Specifies the name of a valid SMS management class to be used to redefine a VSAM cluster and its alternate indexes if any or an IAM file. Your installation ACS routines may override this value.

NEWRECSDATA= Specifies a new primary and optional secondary space in records to be used
NEWRECSD for the data component of a VSAM KSDS or an IAM file, if it is redefined.

Examples:

```
NEWRECSDATA=( 10000 , 1000 )
```

The redefined data component or IAM file will have 10000 records of primary space and 1000 records of secondary space.

```
NEWRECSDATA=10000
```

The redefined data component or IAM file will have 10000 records of primary space and no secondary space.

Note – VSAM will convert the number of records specified into an appropriate number of cylinders or tracks.

NEWRECSINDEX= Specifies a new primary and optional secondary space in records to be used
NEWRECSI for the index component of a VSAM KSDS if it is redefined.

Examples:

```
NEWRECSINDEX=( 500 , 50 )
```

The redefined index component will have 500 records of primary space and 50 records of secondary space.

```
NEWCYLSINDEX=500
```

The redefined index component will have 500 records of primary space and no secondary space.

Note – VSAM will convert the number of records specified into an appropriate number of cylinders or tracks.

NEWSTORCLASS= Specifies the name of a valid SMS storage class to be used to redefine a VSAM cluster and its alternate indexes if any or an IAM file. Your installation ACS routines may override this value.

30.06 CONTINUED . . .

NEWTRKSDATA= Specifies a new primary and optional secondary space in tracks to be used
NEWTRKSD for the data component of a VSAM KSDS or an IAM file, if it is redefined.

Examples:

```
NEWTRKSDATA=( 100 , 10 )
```

The redefined data component or IAM file will have 100 tracks of primary space and 10 tracks of secondary space.

```
NEWTRKSDATA=100
```

The redefined data component or IAM file will have 100 tracks of primary space and no secondary space.

NEWTRKSINDEX= Specifies a new primary and optional secondary space in tracks to be used
NEWTRKSI for the index component of a VSAM KSDS if it is redefined.

Examples:

```
NEWTRKSINDEX=( 2 , 1 )
```

The redefined index component will have 2 tracks of primary space and 1 track of secondary space.

```
NEWTRKSINDEX=2
```

The redefined index component will have 2 tracks of primary space and no secondary space.

NEWVOLSDATA= Specifies up to 20 volumes to be used for the data component of a VSAM
NEWVOLSD KSDS or an IAM file, if it is redefined. If the data set is SMS managed, SMS=SMS will be forced for the define. SMS may or may not use the volumes specified.

Examples:

```
NEWVOLSDATA=(MYVOL 1 , MYVOL 2 )
```

NEWVOLSINDEX= Specifies up to 10 volumes to be used for the index component of a VSAM
NEWVOLSI KSDS if it is redefined. If the data set is SMS managed, SMS=SMS will be forced for the define. SMS may or may not use the volumes specified.

Examples:

```
NEWVOLSINDEX=(MYVOL 1 , MYVOL 2 )
```

NOCKPT Specifies that the checkpoint file should be ignored. This keyword can not be specified with NOLOG. This keyword is required if the JOBNAME keyword is not used and there is no REORGCKP DD statement in the JCL.

NOLOG Specifies that the log file should be ignored. This keyword can not be specified with NOCKPT. This keyword is required if the JOBNAME keyword is not used and there is no REORGLOG DD statement in the JCL.

30.06 CONTINUED . . .

NONSMS=	<p>Specifies if FDRREORG's internal define should keep the current management status for non-SMS managed clusters or let the SMS ACS routines decide.</p> <p>KEEP – NONSMS managed clusters are redefined as non-SMS managed clusters.</p> <p>SMS – The installation SMS ACS routines will decide if the cluster should be defined as managed cluster.</p>
SELTERR=	<p>Specifies if FDRREORG should set a return code of 8 if no data sets are selected by RECOVER.</p> <p>YES – Set return code 8 if no data sets are selected</p> <p>NO – Do not set return code 8 if no data sets are selected.</p> <p>The default is taken from the FDRREORG option table which is shipped set to YES.</p>
SMS=	<p>Specifies if FDRREORG's internal define should keep the current management status and volumes for SMS managed data sets or let the SMS ACS routines decide.</p> <p>KEEP – SMS managed data sets are redefined as SMS managed data sets on the same volumes. On OS/390 systems, only the first volume will be the same.</p> <p>SMS – The installation SMS ACS routines will decide if the data set should be defined as a managed data set. SMS will select the volumes used to define the data set.</p> <p>Note – If the SMS ACS routines indicate that a currently SMS managed data set should not be redefined as SMS managed, FDRREORG will force SMS=KEEP if NEWVOLSDATA and NEWVOLSINDEX were not specified.</p>
SMSADDVOL	<p>Directs FDRREORG to dynamically add an SMS dummy candidate volume for SMS managed data sets before starting the reload. If more space is required during the reload than is available on the current volume(s), the candidate volume will be converted to a real volume and the reload should be able to complete without an out of space error. If the volume is not used, FDRREORG will remove it.</p> <p>Note – This feature is not supported for IAM data sets defined with guaranteed space.</p> <p>Warning – <i>If this keyword was specified on the failed reorganization, another candidate volume will be added during recover processing. If neither candidate is used, only one will be removed.</i></p>

30.06 CONTINUED . . .

SORTRPT=

Specifies the sorting option for the report written to REORGRPT. Specifying YES requires that information related to each data set processed or selected to be kept in storage until the function is completed.

YES – Report is produced at the end of a REORG, SIMULATE or RECOVER function in data set name and volser order.

NO – Report is produced as data sets are processed or selected.

The default is taken from the FDRREORG option table which is shipped set to NO.

VSAMDISP=

Specifies the disposition REORG is to use when allocating VSAM data sets to be recovered.

OLD – VSAM data sets will be allocated DISP=OLD.

SHR – VSAM data sets will be allocated DISP=SHR.

The default and recommended value is OLD !

30.07 FDRREORG OPERATOR COMMANDS

FDRREORG gives the console operator the capability to display the current status of FDRREORG and to stop REORG processing after all active data sets are processed.

**OPERATOR
COMMANDS**

P jobname	FDRREORG will complete processing all active data sets and stop.
F jobname,STATUS	FDRREORG will display on the operators console the volume and data set being processed for each active subtask.

30.10 FDRREORG SIMULATE EXAMPLES

These examples illustrate some of the FDRREORG options. JOB and JOBLIB/STEPLIB DD statements are not shown and if required must specify the load module library in which FDRREORG resides.

**SIMULATE REORG
USING DEFAULT
VALUES**

The following example will report on all VSAM, IAM and PDS files requiring reorganization based on default selection values. (See Section 31 for defaults). All data sets on volume group TEST will be evaluated.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
SIM
SELECT          ALLDSN,VOL=TEST*
/*
```

**SIMULATE REORG
VSAM IF BOTH CI/CA
SPLITS EXCEED 10%**

The following example will generate reports of all VSAM data sets that have a high level qualifier of CICSPROD or PROD, and are on volumes that have a volser beginning with CICS or PROD, and that have split more than 10% of the control area's, and more than 10% of the control intervals.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
SIM
SELECT          DSN=(CICSPROD.*,PROD.*),VOL=(CICS*,PROD*),
                DSTYPE=VSAM,
                IFALL,
                CASPLITR>10,CISPLITR>10
/*
```

**SIMULATE REORG
VSAM IF EITHER THE
CA SPLITS EXCEED 5%
or CI SPLITS EXCEED
10%**

The following example will generate reports of all VSAM data sets that have a high level qualifier of APPLIC from the catalog, and that have split either more than 5% of the control area's, or more than 10% of the control intervals.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
SIM
SELECT          CATDSN=APPLIC.*,
                DSTYPE=VSAM,
                IFANY,
                CASPLITR>5,CISPLITR>10
/*
```

30.10 CONTINUED . . .

**SIMULATE
REORG OF IAM
FILES IF EITHER
PE OR
OVERFLOW
EXCEED 75%**

The following example will generate reports of all IAM data sets that have \$IAM anywhere in the data set name, and that have used more than 75% of the overflow area or have used more than 75% of the prime extension.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
SIM
SELECT          CATDSN=**$IAM**,
                DSTYPE=IAM,IFANY,
                OFULL>75,PEFULL>75

/*
```

**SIMULATE
PDS's LARGER
THAN ONE
TRACK**

The following example will generate reports of all partitioned data sets that are larger than 1 track and are at least 90% full or are in more than 5 extents that reside on volumes that have a VOLSER beginning with TSO.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
SIM
SELECT          ALLDSN,VOL=TSO*,
                DSTYPE=PDS,
                TRKS>1,
                IFANY,
                PDSFULL>=90,PDSEXTENTS>5

/*
```

**SIMULATE
PDS's USING
FDRCOPY**

The following example will report on the number of tracks that would be reclaimed by compressing PDS's with an ending qualifier of CNTL or LOAD on any online volume.

Note: PDS compression can be done with FDRREORG or FDRCOPY. SIMREORG requires that FDR read all used tracks in the selected PDS's.

```
//STEPNAME      EXEC   PGM=FDRCOPY,REGION=2048K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//SYSIN          DD     *
SIMREORG        TYPE=DSF,DSNENQ=TEST
SELECT          DSN=**.*CNTL,VOL=*
SELECT          DSN=**.*LOAD,VOL=*

/*
```

30.10 CONTINUED . . .

**SIMULATE
COMBINED
DATA SET
TYPES**

The following example shows how to select different criteria for different types of data sets. VSAM files will be selected if CA or CI splits exceed 5%. IAM files will be selected if Independent Overflow exceeds 50% and PDS data sets will be selected if they exceed 1 track and are over 80% full.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
SIM
  SELECT         VOL=(CICS*,PROD*),
                 DSN=(CICSPROD.***,PROD.**),
                 DSTYPE=VSAM,
                 IFANY,
                 CASPLITR>5,CISPLITR>5
  SELECT         VOL=(CICS*,PROD*),
                 DSN=(CICSPROD.***,PROD.**),
                 DSTYPE=IAM,
                 OFULL>50
  SELECT         DSN=**,VOL=TSO*,
                 DSTYPE=PDS,
                 TRKS>1,
                 PDSFULL>=80
/*
```

30.11 FDRREORG EXAMPLES

**REORG USING
TEMPORARY
BACKUPS**

The following example can be used to reorganize all CICSPROD. and PROD. data sets on any volume with a volser beginning with CICS or PROD that have split more than 10% of the control areas or more than 20% of the control intervals. The name of the backup data sets will be the name of the selected data set appended with .BACKUP. Each backup data set will be deleted after the selected data set is reorganized.

```
//STEPNAME EXEC PGM=FDRREORG,REGION=4096K
//SYSPRINT DD SYSOUT=*
//REORGPRPT DD SYSOUT=*
//REOGRPT DD SYSOUT=*
//SYSIN DD *
REORG BACKUP=TEMP,BACKUPUNIT=SYSDA,
      BACKUPINDEX=++BACKUP
      SELECT VOL=(CICS*,PROD*),
            DSN=(CICSPROD.***,PROD.***),
            DSTYPE=VSAM,
            IFANY,
            CASPLITR>10,CISPLITR>20
```

**REORG USING
PERMANENT
BACKUPS**

In this run, BACKUP=PERM is used to keep the backup data sets. If FDRREORG is run at a later date with the same control cards, any backup data sets that were kept from a previous run must be deleted first or the data set will not be reorganized. FDRREORG will not create a backup using the name of an existing cataloged data set.

```
//STEPNAME EXEC PGM=FDRREORG,REGION=4096K
//SYSPRINT DD SYSOUT=*
//REORGPRPT DD SYSOUT=*
//REOGRPT DD SYSOUT=*
//SYSIN DD *
REORG BACKUP=PERM,BACKUPUNIT=SYSDA,
      BACKUPINDEX=++BACKUP
      SELECT VOL=(CICS*,PROD*),
            DSN=(CICSPROD.***,PROD.***),
            DSTYPE=VSAM,
            CASPLITR>10,CISPLITR>20
```

**REORG USING
-
ALWAYSBACK
UP**

To combine an application level backup with reorganization, the ALWAYSBACKUP option can be used. With this option, all data sets that pass the basic selection criteria will be backed up, and any data sets that meet the reorganization criteria will also be reorganized. Use of the ALWAYSBACKUP option requires that the backup data sets be kept. BACKUP=TEMP, if specified or defaulted, will be treated as a keyword error. In the following example, all CICSPROD and PROD PAYROLL VSAM KSDSs or AIXs will be backed up. If any data sets are encountered that have a CI or CA split ratio greater than 10, will be reorganized after the backup.

```
//STEPNAME EXEC PGM=FDRREORG,REGION=
//SYSPRINT DD SYSOUT=*
//REORGPRPT DD SYSOUT=*
//REOGRPT DD SYSOUT=*
//SYSIN DD *
REORG BACKUP=GDG,BACKUPUNIT=tape,
      BACKUPINDEX=++BACKUP
      SELECT VOL=(CICS*,PROD*),
            DSN=(CICSPROD.PAYROLL.***,PROD.PAYROLL.***),
            ALWAYSBACKUP,
            DSTYPE=VSAM,
            IFANY,
            CASPLITR>10,CISPLITR>10
/*
```

30.11 CONTINUED . . .

**REORG USING
GDG BACKUPS
ON tape**

The following example will select data sets with high level index of ACCOUNT from the catalog.

In this run, BACKUP=GDG is used to keep the backup data sets as generation data sets. Using this method allows the data sets to be kept without having to delete old backups first. FDRREORG will dynamically define the GDG base record if one does not already exist.

The backup files will use tape.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REOGRPT        DD     SYSOUT=*
//SYSIN          DD     *
REORG            BACKUP=GDG,BACKUPUNIT=3480,
                  BACKUPINDEX=++BACKUP
SELECT           CATDSN=ACCOUNT.**,
                  DSTYPE=VSAM,
                  IFANY,
                  CASPLITR>5,CISPLITR>10
```

**REORG
SPECIFIC
DATA SETS**

The following example demonstrates how to select specific data sets for reorganization whether they require reorganization or not. All VSAM, IAM or PDS data sets with a second index level of ACCOUNT will be selected from volumes ACCT01, ACCT09 and ACCT10.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REOGRPT        DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS
SELECT           DSN=*.ACCOUNT.**,
                  VOL=(ACCT01,ACCT09,ACCT10)
```

**REORG VSAM
IF ANY SPLITS
OCCUR**

The following examples will reorganize any VSAM file with an ending index level of cluster and have experienced any CI or CA splits.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REOGRPT        DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS
SELECT           DSN=**.CLUSTER,VOL=PAY*,
                  DSTYPE=VSAM,
                  NCISPLITS>0,NCASPLITS>0
```

**REORG USING
- NOREORG**

FDRREORG can be used to take application level backups by using the NOREORG option. With this option, all data sets that meet the selection criteria will be backed up but not reorganized. Use of the NOREORG option requires that the backup data sets be kept. BACKUP=TEMP, if specified or defaulted, will be treated as a keyword error. In the following example, all CICS PROD and PROD PAYROLL VSAM KSDSs or AIXs will be backed up.

NOTE: The RECOVER statement is not designed to restore these backups. You can use an IDCAMS REPRO JOB to reload these data sets.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=...
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REOGRPT        DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS,BACKUP=GDG,BACKUPUNIT=tape,
                  BACKUPINDEX=++BACKUP
SELECT           VOL=(CICS*,PROD*),
                  DSN=(CICSPROD.PAYROLL.**,PROD.PAYROLL.**),
                  DSTYPE=VSAM,
                  NOREORG
```

30.11 CONTINUED . . .

**REORG USING
AN EXCLUDE**

The following example will select data sets starting with CICSPROD or PROD with the exception of any data sets that have a second level qualifier of PAYROLL. Note that control cards are processed in the order you enter them. If the EXCLUDE statement was placed after the SELECT statement, the CICSPROD.PAYROLL and PROD.PAYROLL data sets would be reorganized because they would be selected before the EXCLUDE statement is processed. The backup data set name will use the RACF userid as its high level index and will go to SMS STORCLAS of TEMPdisk.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPRT      DD     SYSOUT=*
//SYSIN          DD     *
  REORG          BACKUP=TEMP,BACKUPSTORCLASS=TEMPdisk,
                  BACKUPINDEX=&RACFUID.++BACKUP
  EXCLUDE        DSN=(CICSPROD.PAYROLL.***,PROD.PAYROLL.***)
  SELECT         VOL=(CICS*,PROD*),
                  DSN=(CICSPROD.***,PROD.***),
                  DSTYPE=VSAM,
                  IFANY,
                  CASPLITR>10,CISPLITR>20
/*
```

**REORG USING
CRDAYS**

The following example will select VSAM files with an additional filter limiting the selection to data sets created within the last two days. This might be useful if new data sets were defined with incorrect free space parameters that have been modified to the correct values with an IDCAMS ALTER statement. By reorganizing these data sets, they will be reloaded using the new free space parameters.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPRT      DD     SYSOUT=*
//SYSIN          DD     *
  REORG          BACKUP=TEMP,BACKUPUNIT=SYSDA,
                  BACKUPINDEX=++BACKUP
  SELECT         VOL=(CICS*,PROD*),
                  DSN=(CICSPROD.***,PROD.***),
                  DSTYPE=VSAM,
                  CRDAYS<=2          <---select if created in last 2 days
/*
```

You could also achieve the same results by using specific dates as shown in the following example. This example also demonstrates how to specify a low and high range for selection.

```
  REORG          BACKUP=TEMP,BACKUPUNIT=SYSDA,
                  BACKUPINDEX=++BACKUP
  SELECT         VOL=(CICS*,PROD*),
                  DSN=(CICSPROD.***,PROD.***),
                  DSTYPE=VSAM,
                  CRDATE>=98001,CRDATE<=98002 <---date range
/*
```

or

```
  CRDATE>=1998001,CRDATE<=1998002
/*
```

30.11 CONTINUED . . .

**REORG USING
LAST
REFERENCE
DATE AS
FILTERS**

The last reference date can also be used as a filter. Use LRDATE if you want to filter on a specific date. Use LRDAY5 to filter based on days since last reference.

The following example will select partitioned data sets that have been referenced within the last 7 days and are more than 80% full. This allows FDR to bypass PDSs that have not been referenced in some time and most likely do not require compression.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=6M
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRT      DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
REORG
SELECT          ALLDSN,VOL=TSO*,
                DSTYPE=PDS,
                LRDAY5<=7,  <---referenced in last 7 days
                PDSFULL>80  <---must be more than 80% full

/*
```

**COMPRESS
PDS's USING
FDRCOPY FOR
PERFORMANCE**

The following example uses FDRCOPY to compress a large number of PDS data sets. FDRCOPY can be used instead of FDRREORG to compress PDS data sets if the special features of FDRREORG (ex: extra selection criteria, etc.) are not required. FDRCOPY reduces the overhead of compressing large number of data sets yielding optimum performance.

```
//STEPNAME      EXEC   PGM=FDRCOPY,REGION=2M
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//SYSIN         DD     *
REORG          TYPE=DSF
SELECT         ALLDSN,VOL=*

/*
```

**REORG LARGE
FILES TO TAPE
SMALL TO
DISK**

The following example will reorganize VSAM and IAM files based on default selection criteria. Files 100 cylinders or larger will use tape as the backup medium. Files under 100 cylinders will use disk for the backups. Since a large number of volumes will be processed 8 subtasks will be used.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=8M
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRT      DD     SYSOUT=*
//REORGRPT      DD     *
//SYSIN         DD     *
REORG          DATA=USED,MAXTASKS=8,
                BACKUPUNIT=DISK
SELECT         ALLDSN,VOL=MANY*,
                DSTYPE=(VSAM,IAM),
                CYLS<100
REORG          DATA=USED,MAXTASKS=8,
                BACKUPUNIT=TAPE
SELECT         ALLDSN,VOL=MANY*,
                DSTYPE=(VSAM,IAM),
                CYLS>=100

/*
```


30.11 CONTINUED . . .

**REORG
EXAMPLES -
UNAVAILABLE
DATA SETS**

If data sets are in use at the time FDRREORG selects them for processing, they will normally be bypassed. In some cases it might be critical that certain data sets be reorganized before starting an on-line region, or before starting a production job stream. In other cases the data set might become available before FDRREORG processing is complete. These situations can be handled by using the DSNRETRY parameter. For data sets that you would like processed if they become available, you can use DSNRETRY=RETRY or DSNRETRY=ENQ. Specifying RETRY instructs FDRREORG to make repeated attempts to allocate the data set during the normal course of processing. If the data set is successfully allocated, it will be processed. ENQ instructs FDRREORG to leave pending a SYSDSN ENQ for the data set. This ensures that FDRREORG will obtain ownership of the data set as soon as it becomes available. If the data set ENQ is not obtained before the task that selected the data set completes, the pending SYSDSN ENQ will be released and the data set will not be processed. For those situations where it is critical that data sets be reorganized, you can specify DSNRETRY=WAIT. In this case, the task that selected the data set will wait for the SYSDSN ENQ to be obtained and will not terminate until all data sets selected with DSNRETRY=WAIT have been processed. We have added RUNTIME=60 to the REORG statement which instructs FDRREORG to stop processing after 60 minutes. If RUNTIME was not specified on the REORG statement, or an operator STOP (P) command would have to be issued in the event that a data set queued for retry processing with the wait option will be unavailable for a long period of time.

The following example demonstrate how to use the various DSNRETRY options. The first SELECT statement causes each task to wait for CICSProd. data sets on CICSxx volumes that are not available when selected. The second SELECT statement causes each task to issue a SYSDSN ENQ for PROD. data sets on PRODxx volumes that are 1000 cylinders or smaller. These data sets will not be processed if they do not become available before the task that selected the data sets completes its processing. The third SELECT statement causes each task to make repeated attempts to allocate PROD. data sets on PRODxx volumes that are larger than 1000 cylinders. A SYSDSN ENQ will not be issued for these data sets and they will not be processed if all allocation attempts are unsuccessful.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRTPRT    DD     SYSOUT=*
//REORGPRTPRT    DD     SYSOUT=*
//SYSIN          DD     *
  REORG          BACKUP=TEMP,BACKUPUNIT=SYSDA,
                  BACKUPINDEX=++BACKUP,
                  DATA=USED,
                  DSTYPE=VSAM,
                  RUNTIME=60
  SELECT         VOL=(CICS*),
                  DSN=(CICSProd.**),
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=WAIT
  SELECT         VOL=(PROD*),
                  DSN=(PROD.**),
                  CYLS<=1000,
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=ENQ
  SELECT         VOL=(PROD*),
                  DSN=(PROD.**),
                  CYLS>1000,
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=RETRY
/*
```

30.11 CONTINUED . . .

**CHANGING
THE SPACE
ALLOCATION
OF A VSAM
CLUSTER**

In the following example, the physical allocation of a VSAM cluster will be changed. The data component will be allocated with 100 cylinders primary and 20 cylinders secondary. The index component will be allocated with 5 tracks primary and 1 track secondary.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
      REORG       NODEFAULTS,NOUPDATES=YES
      SELECT      CATDSN=VSAM.CLUSTER,
                  NEWCYLSDATA=(100,20),NEWTRKSINDEX=(5,1)

/*
```

**MOVING
VSAM DATA
SETS**

In the following example, the data and index components of a VSAM cluster will be redefined on volume VSAM01.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
      REORG       NODEFAULTS,NOUPDATES=YES
      SELECT      NEWVOLSDATA=VSAM01,NEWVOLINDEX=VSAM01

/*
```

**MOVING
VSAM DATA
SETS WITH
ALTERNATE
INDEXES**

In the following example, the data and index components of a VSAM cluster residing on a 3380 will be redefined on volume VSAM01 (a 3390 volume). By specifying MOVEAIX=YES, any alternate indexes of this VSAM cluster will also be redefined on volume VSAM01. If MOVEAIX=YES is not specified, the default is NO and the alternate indexes will be defined on their original volumes.

```
//JOBNAME        JOB    .....
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
      REORG       NODEFAULTS,NOUPDATES=YES,MOVEAIX=YES
      SELECT      CATDSN=VSAM.CLUSTER,
                  NEWVOLSDATA=VSAM01,NEWVOLINDEX=VSAM01

/*
```

**CONVERT A
NON-SMS
VSAM
CLUSTER TO
AN SMS
MANAGED
CLUSTER**

In the following example, a VSAM cluster will be redefined as an SMS managed cluster. Please note that the installation ACS routines can override the requested storage class or indicate that the cluster should be defined as a NON-SMS managed cluster.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
      REORG       NODEFAULTS,NOUPDATES=YES
      SELECT      CATDSN=VSAM.CLUSTER,NEWSTORCLASS=PROD

/*
```

30.11 CONTINUED . . .

**CHANGING
THE SPACE
ALLOCATION
OF AN IAM
DATASET**

In the following example, the physical allocation of an IAM data set will be changed. The file will be allocated with 100 cylinders primary and 20 cylinders secondary.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=0M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRT       DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS,NOUPDATES=YES
SELECT           CATDSN=IAM.DATASET,NEWCYLSDATA=(100,20)
/*
```

**MOVING IAM
DATASETS**

In the following example, an IAM data set will be redefined on volume IAM001.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=0M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRT       DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS,NOUPDATES=YES
SELECT           CATDSN=IAM.DATASET,NEWVOLSDATA=IAM001
/*
```

**CONVERT A
NON-SMS IAM
DATASET TO
AN SMS
MANAGED
DATASET**

In the following example, an IAM data set will be redefined as an SMS managed data set. Please note that the installation ACS routines can override the requested storage class or indicate that the data set should be defined as a NON-SMS managed data set.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=0M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRT       DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
REORG            NODEFAULTS,NOUPDATES=YES
SELECT           CATDSN=IAM.DATASET,NEWSTORCLASS=PROD
/*
```

**PARALLEL
MODE
REORGANIZATION**

In the following example, FDRREORG is run in parallel mode for a multi-volume IAM and VSAM file. A maximum of 9 parallel tasks per data set are allowed as specified by the MAXP keyword. Note that if the two data sets are using space on 9 or more volumes, 18 tape drives will be allocated assuming that the first volume of each data set are not the same.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=0M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRT       DD     SYSOUT=*
//REORGRPT       DD     SYSOUT=*
//SYSIN          DD     *
REORG            MODE=P,MAXP=9,BACKUPUNIT=3480,MAXT=2,
                  BACKUPINDEX=++BACKUP?
SELECT           CATDSN=(VSAM.CLUSTER,IAM.DATASET)
```

30.11 CONTINUED . . .

**COMBINED
PARALLEL
AND SINGLE
MODE
REORGANIZATION**

In the following example, FDRREORG is first run in parallel mode for multi-volume IAM and VSAM files. After all of the parallel mode reorganizations are complete, the second REORG and SELECT statements will reorganize all of the data sets that were not reorganized in parallel mode. A maximum of 5 parallel tasks per data set are allowed as specified by the MAXP keyword. The MAXTASKS keyword has not been specified which will cause FDRREORG to use the default value of 4 if the default has not been changed in the FDRREORG option table. This means that the parallel portion of the reorganization could require 20 tape drives, 4 data sets times 5 parallel backups per data set. To use less tape drives, specify smaller values for MAXP and MAXTASKS.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=0M
//SYSUDUMP       DD     SYSOUT=*
//SYSPRINT       DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//REORGPRPT      DD     SYSOUT=*
//SYSIN          DD     *
REORG            MODE=P,MAXP=5,BACKUPUNIT=3480,BACKUPINDEX=++BACKUP?,
                  DSTYPE=(IAM,VSAM)
SELECT           CATDSN=(PROD.***)
REORG            MODE=S,BACKUPUNIT=3480,BACKUPINDEX=++BACKUP,
                  DSTYPE=(IAM,VSAM)
SELECT           CATDSN=(PROD.***)
```

30.11 CONTINUED . . .

**REORG EXAMPLES –
UNAVAILABLE
DATASETS**

If data sets are in use at the time FDRREORG selects them for processing, they will normally be bypassed. In some cases it might be critical that certain data sets be reorganized before starting an on-line region, or before starting a production job stream. In other cases the data set might become available before FDRREORG processing is complete. These situations can be handled by using the DSNRETRY parameter. For data sets that you would like processed if they become available, you can use DSNRETRY=RETRY or DSNRETRY=ENQ. Specifying RETRY instructs FDRREORG to make repeated attempts to allocate the data set during the normal course of processing. If the data set is successfully allocated, it will be processed. ENQ instructs FDRREORG to leave pending a SYSDSN ENQ for the data set. This ensures that FDRREORG will obtain ownership of the data set as soon as it becomes available. If the data set ENQ is not obtained before the task that selected the data set completes, the pending SYSDSN ENQ will be released and the data set will not be processed. For those situations where it is critical that data sets be reorganized, you can specify DSNRETRY=WAIT. In this case, the task that selected the data set will wait for the SYSDSN ENQ to be obtained and will not terminate until all data sets selected with DSNRETRY=WAIT have been processed. We have added RUNTIME=60 to the REORG statement which instructs FDRREORG to stop processing after 60 minutes. If RUNTIME was not specified on the REORG statement, or an operator STOP (P) command would have to be issued in the event that a data set queued for retry processing with the wait option will be unavailable for a long period of time.

The following example demonstrates how to use the various DSNRETRY options. The first SELECT statement causes each task to wait for CICS PROD. Data sets on CICSxx volumes that are not available when selected. The second SELECT statement causes each task to issue a SYSDSN ENQ for PROD. Data sets on PRODxx volumes that are 1000 cylinders or smaller. These data sets will not be processed if they do not become available before the task that selected the data sets completes its processing. The third SELECT statement causes each task to make repeated attempts to allocate PROD. Data sets on PRODxx volumes that are larger than 1000 cylinders. A SYSDSN ENQ will not be issued for these data sets and they will not be processed if all allocation attempts are unsuccessful.

```
//STEPNAME      EXEC    PGM=FDRREORG,REGION=4096K
//SYSUDUMP       DD      SYSOUT=*
//SYSPRINT       DD      SYSOUT=*
//REORGPRT      DD      SYSOUT=*
//REORGPRPT     DD      SYSOUT=*
//SYSIN         DD      *
  REORG          BACKUP=TEMP,BACKUPUNIT=SYSDA,
                  BACKUPINDEX=++BACKUP,
                  DATA=USED,
                  DSTYPE=VSAM,
                  RUNTIME=60
  SELECT        VOL=(CICS*),
                  DSN=(CICS PROD.**),
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=WAIT
  SELECT        VOL=(PROD*),
                  DSN=(PROD.**),
                  CYLS<=1000,
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=ENQ
  SELECT        VOL=(PROD*),
                  DSN=(PROD.**),
                  CYLS>1000,
                  IFANY,CASPLITR>10,CISPLITR>20,
                  DSNRETRY=RETRY
/*
```

30.12 FDRREORG RECOVER EXAMPLES

RECOVERY EXAMPLE The following sample control card demonstrates how you can recover from a failed reorganization. This example assumes that the jobname of the failed reorganization job was MYREORG.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
  RECOVER       JOBNAME=MYREORG
/*
```

The RECOVER statement assumes that you want to complete reload processing for all data sets found in the checkpoint or log files. You can use the DSN keyword to provide a data set or data set filter list to limit recovery processing to those data sets you provide, or you can use the EXDSN keyword to exclude data sets from recovery processing. The following sample control cards demonstrate how you can use the DSN or EXDSN keywords. In the first example, only data sets with a high level index of PROD will be recovered. In the second example, only data sets that do not have a high level index of PROD will be recovered.

```
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
  RECOVER       JOBNAME=MYREORG,DSN=PROD.**
/*
//STEPNAME      EXEC   PGM=FDRREORG,REGION=4096K
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//REORGPRPT     DD     SYSOUT=*
//REORGRPT      DD     SYSOUT=*
//SYSIN         DD     *
  RECOVER       JOBNAME=MYREORG,EXDSN=PROD.**
/*
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

31.01 FDRREORG OPTIONS AND DEFAULTS

FDRREORG has its own option table, separate from the FDR Global Option Table. Like the FDR Option Table, it is stored as a load module in the FDR program library, under the name FDRREOOP. This module is loaded every time FDRREORG is executed to query the options and defaults that have been set. The installation options include default selection criteria, special features in FDRREORG, and other processing options.

FDRREORG also has a NOREORG list, used to specify data sets which should never be reorganized. It is also stored in the FDR program library under the name FDRNORG.

SETTING OPTIONS AND DEFAULTS

There are 3 ways to change the settings in the FDRREORG option table or update the NOREORG list:

- you can use the FDR ISPF installation panels, as described in [Section 90.23](#) and [90.31](#). **This is the recommended method**
- you can execute the FDRREORG Option Change Utility, FDRREOZO, in a batch job step
- you can execute FDRREOZO under TSO.

DEFAULTS

FDRREORG is shipped with default values set in the FDRREORG option table for selected keywords if you do not provide your own. If you code a REORG statement and do not provide the keywords or an alternative to the keywords shown below, the REORG statement will be executed as if you provided the following control card:

```
REORG          BACKUP=TEMP, BACKUPINDEX=++BACKUP,
                BACKUPALLOC=UNIT, BACKUPUNIT=SYSDA,
                CKPTPREFIX=FDRREORG, CKPTALLOC=UNIT, CKPTUNIT=SYSDA,
                LOGPREFIX=FDRREORG, LOGALLOC=UNIT, LOGUNIT=SYSDA,
                DATA=USED, DSNRETRY=NO,
                OWNERSTRING=$FDR, VSAMDEFINE=IFREQ, BUILDEMPTYAIX=NO
```

For the SELECT statement, you must provide at least one and it must contain a data set or data set filter list specified by the CATDSN or DSN keywords, and a volume list specified by the VOL keyword if the DSN keyword was used to provide the data set or data set filter list. If you do not provide the keywords or an alternative to the keywords shown below, the SELECT statement is evaluated as if you provided the following control card:

```
SELECT  CATDSN=yourlist,          <---user must provide
        DSN=yourlist, VOL=yourlist, or
        DSTYPE=ALL, FREESPACE<50, <---user must provide
        IFANY,                   <---VSAM default
        CASPLITR>=10, CISPLITR>=10, <---VSAM defaults
        PDSFULL>=90, PDSEXTENTS>=17, <---PDS defaults
        OFULL>=90, PEFULL>=100,    <---IAM defaults (compatible format only)
        PCTRECO>=10, PEUDATAR>=100 <---IAM defaults (all formats)
        OVERFLOWINDEX=1048576    <---IAM defaults (all formats)
```

If these keywords are not specified on a SELECT statement, the default values will be used for the keywords not specified. If you want the SELECT statement to be evaluated using only the values specified, code the NODEFAULTS keyword.

31.01 FDRREOZO JCL REQUIREMENTS

EXEC STATEMENT	Must specify the program name of the FDRREORG Option Change Facility – FDRREOZO.
STEPLIB or JOBLIB DD STATEMENT	If required, must specify the load module library in which FDRREORG resides.
SYSPRINT DD STATEMENT	Specifies the output message data set. This is a required DD statement and usually is a SYSOUT data set.
SYSLIB DD STATEMENT	Must specify the load module library in which FDRREORG resides.
SYSIN DD STATEMENT	Specifies the control statement data set required for all functions. Usually an input stream or DD * data set.
TSO EXECUTION	<p>FDRREOZO can be executed under TSO. The program will prompt the user for the commands. 'END' will save the new options, if any, and terminate the program. The allocations required to execute FDRREOZO in the TSO Foreground are as follows:</p> <pre> ALLOC F(SYSLIB) DA('fdrreorg.library') SHR ALLOC F(SYSPRINT) DA(*) ALLOC F(SYSIN) DA(*) FDRREOZO ---- or if the FDRREORG library is not in LINKLIST ---- CALL 'fdrreorg.library(FDRREOZO)'</pre>
ISPF PANELS	All of the FDRREORG options can be set through the FDR ISPF installation panels as described in Sections 90.23 (FDRREORG options) and 90.31 (NOREORG table). This is the recommended method for updating FDRREORG options.

31.03 FDRREOZO STATEMENTS

**FDRREOZO
STATEMENTS**

FDRREOZO accepts these statements:

- CANCEL**— Terminates FDRREOZO without updating the Options Table.
- END** — Terminates FDRREOZO processing and rewrites the Option Table if any option was changed. This command is required for TSO users.
- NOREORG**— Adds entries to the user modifiable section of the NOREORG list which is located in the FDRNORG load module. It is documented in [Section 31.04](#).
- PRINT** — Display the current values in the FDRREORG Options table and/or the NOREORG list.
The PRINT statement format is:
**PRINT or
PRINT NOREORG or
PRINT ALL**
- RESET** — Reset the FDRREORG Options to the original values on the installation tape and/or reset the NOREORG list by removing all entries from the user modifiable section of the NOREORG list.
The RESET statement format is:
**RESET or
RESET NOREORG or
RESET ALL**
- ZAP** — Modify the FDRREORG Options. It is documented in [Section 31.05](#).

31.04 FDRREOZO NOREORG STATEMENTS

NOREORG DSN=(dsn1,...dsnn),VOL=(vol1,...voln)

NOREORG STATEMENT The NOREORG statement adds entries the user-modifiable part of the FDRREORG NOREORG table. The named data sets and/or volumes will never be reorganized by FDRREORG.

FDRREOZO does not provide a facility for selectively deleting data sets from the table. Innovation suggests that you create and maintain a FDRREOZO jobstream starting with a RESET NOREORG statement to empty the table, followed by NOREORG statements to rebuild the entire table. This allows you to edit the NOREORG statements and rebuild the table at any time.

OPERANDS	DSN=	Specifies a data set name or a data set filter name (described in Section 80.14) which is not to be reorganized. Multiple names or filters can be specified by enclosing the list in parenthesis. If VOL= is also specified, only data sets matching the name or filter on the named volumes are protected; otherwise they are protected if found on any volume.
	VOL=	Specifies a disk volume serial or a volser prefix (followed by an asterisk, e.g., VOL=ABC*) which is not to be reorganized. Multiple serials or prefixes can be specified by enclosing the list in parenthesis. If DSN= is also specified, only data sets matching the name or filter on the named volumes are protected; otherwise all data sets on the named volumes are protected.

31.05 FDRREOZO ZAP STATEMENT

ZAP

Default VSAM
Selection Criteria

CASPLITR=nnnnn

,CISPLITR=nnnnn

,FREESPACE=nnn

Default IAM
Selection Criteria

,OFULL=nnn

,OVERFLOWINDEX=nnnnnnnnnn

,PCTTRECO=nnn

,PEFULL=nnn

,PEUDATAR=nnn

Default PDS
Selection Criteria

,PDSFULL=nnn

,PDSEXTENTS=nn

REORG, SIMULATE And RECOVER Defaults

,ALIASCHECK=YES|NO

,AUTOSPACE=YES|NO

,BACKUP=TEMP|PERM|GDG

,BACKUPALLOC=SMS|UNIT

,BACKUPDATACLASS=dataclas

,BACKUPFAILURE=ERROR|WARNING

,BACKUPINDEX=index

,BACKUPMGMTCLASS=mgmtclas

,BACKUPRETPD=nnn

,BACKUPSTORCLASS=storclas

,BACKUPUNIT=unitname

,BACKUPUNITS=n

,CKPTALLOC=SMS|UNIT

,CKPTDATACLASS=dataclas

,CKPTMGMTCLASS=mgmtclas

,CKPTPREFIX=prefix

,CKPTSTORCLASS=storclas

,CKPTUNIT=unitname

,ENQERR=YES|WRN|NO

,IAMCOMPPERM=YES|NO

,IAMCOMPTEMP=YES|NO

,IAMDEFINE=YES|NO

,KEEPCRDATE=YES|NO

,LASTAPEPREFIX=prefix

,LISTNOREORG=YES|NO

,LOGALLOC=SMS|UNIT

,LOGDATACLASS=dataclas

,LOGMGMTCLASS=mgmtclas

,LOGPREFIX=prefix

,LOGSTORCLASS=storclas

,LOGUNIT=unitname

,MAXENQ=nnn

,MAXENVERR=nnnnn

,MAXFILE=nnn

,MAXSYSERR=nnnnn

,MAXTASKS=nn

,MSGLEVEL=c

,MSGTIMESTAMP=ccc

,NOUPDATES=ccc

,OWNERSTRING=cccc

,POOLDASD=ccc

,RECORDCOUNTS=ccc

,REDUCEPRIMARY=

,SELTERR=ccc

,SORTRPT=ccc

,SUBTASKABEND=cccc

,UPDATEDPDS=ccc

,VSAMDEFINE=cccccc

CONTINUED . . .

31.05 CONTINUED . . .

ZAP STATEMENT The ZAP statement is used to modify individual FDRREORG options and defaults in the FDRREORG Option Table. The ZAP statement may contain one or more of the operands described here; you can also have multiple ZAP statements. Only the options you specify are changed. The modifications are made to a copy of the table in memory; it is not saved back to the FDR program library until the END statement is entered or (for a batch job) the end of SYSIN is encountered. The CANCEL command will discard the changes.

Note: the defaults shown under each operand are the defaults which are distributed with FDRREORG. If you have previously updated some FDRREORG options, they will have the default of whatever you have currently set them to. To reset all options to the defaults distributed by Innovation, use the RESET statement.

31.05 CONTINUED . . .

OPERANDS ALIASCHECK=

Specifies the default option for verifying if a catalog alias exists for generated backup data set names.

YES – Check for an existing alias and fail the backup or reorganization if one does not exist.

NO – Do not check for an existing alias. If an alias does not exist, the backup data set will be cataloged in the master catalog. This will fail if the job does not have authority to update the master catalog.

The default is YES.

AUTOSPACE=

Specifies if FDRREORG should attempt to increase the space allocation of single volume IAM or VSAM data sets if a reload fails due to an out of space condition. If YES is specified, FDRREORG will calculate how much additional space is required to reload the data set. The data set will then be redefined with additional space and the reload will be retried.

YES – Out of space failures will be handled by FDRREORG.

NO – Out of space failures will not be handled by FDRREORG.

The default is YES.

Note – This feature is deactivated for data sets redefined at reload time with space parameters provided via the NEWCYLS, NEWTRKS, or NEWRECS keywords for the data or index components.

For IAM data sets, this feature requires IAM 6.3 or above. Additionally, files defined with an earlier release of IAM will not be supported by this feature until they are first redefined with IAM 6.3 or above.

BACKUP=

Defines the default disposition of the backup data sets after the target data set has been successfully reorganized.

TEMP – Backup data sets on disk will be deleted and uncataloged. Backup data sets on tape will be uncataloged.

PERM – Backup data sets on disk will not be deleted and uncataloged. Backup data sets on tape will not be uncataloged.

GDG – Same as PERM except backup data sets will be a +1 generation data set. REORG will dynamically define any generation data groups that do not already exist. The high level qualifier of the dynamically generated generation data groups must be aliased to an ICF catalog.

The default is TEMP.

31.05 CONTINUED . . .

**BACKUPALLOC=
BACKUPA**

Defines the allocation method to use when allocating backup data sets.

SMS – Allocate by SMS storage class.

UNIT – Allocate by unit name.

Note- If SMS is specified, a valid SMS storage class name must also be specified for BACKUPSTORCLASS. If UNIT is specified, a valid unit name must also be specified for BACKUPUNIT.

The default is UNIT.

**BACKUPDATACLASS=
BACKUPD**

Defines the default SMS data class to use for backup data sets when BACKUPALLOC is SMS.

BACKUPFAILURE=

Defines how FDRREORG should treat a backup failure that occurs as part of a reorganization. Backup failures that occur as a result of data sets selected via the NOREORG or ALWAYSBACKUP keywords will always be treated as an error.

ERROR – The failure is treated as an error and will result in a return code of 8.

WARNING – The failure is treated as a warning and will result in a return code of 4.

The default is ERROR.

**BACKUPINDEX=
BACKUPI**

Defines the default pattern to be used to add or delete index levels when generating the backup data set name. REORG will use each index level specified in BACKUPINDEX in place of the original index level. If a period is specified without any characters following, the original index level will be copied to the backup data set name. If + is specified, the character following the + will be inserted into the backup data set name as a new index level. If ++ is specified, the characters following the ++ will be added to the end of the backup data set name as a new index level. If – is specified, the index level will be dropped from the backup data set name. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used. See the description of the SELECT statement for examples.

The default is ++BACKUP.

**BACKUPMGMTCLASS=
BACKUPM**

Defines the default SMS management class to use for backup data sets when BACKUPALLOC is SMS.

**BACKUPRETPD=
BACKUPR**

Defines the default retention period to be used for backup data sets.

The default is 0.

**BACKUPSTORCLASS=
BACKUPS**

Defines a valid SMS storage class which will be used to allocate backup data sets when BACKUPALLOC is SMS.

BACKUPUNIT=

Defines a valid unit name containing tape or disk devices to be used for allocating backup data sets when BACKUPALLOC is UNIT. The unit name specified must not define a group of devices that contain a mixture of device classes (ie. 3480s and 3380s) . It is allowed to use a unit name that defines a group of devices with a mixture of device models (ie. 3380s and 3390s).

The default is SYSDA.

31.05 CONTINUED . . .

BACKUPUNITS=	<p>For backup data sets on tape, this operand defines the number of units to be allocated for backup data sets for each task. For backup data sets on disk, this operand defines the minimum number of units to be allocated for each backup data set. REORG will dynamically increase the number of disk units for a disk backup data set to ensure that sufficient space is available to backup the target data set.</p> <p>The default is 1.</p>
CASPLITR= CA	<p>Defines the default minimum ratio of Control Area splits to every 100 Control Area's in a VSAM KSDS.</p> <p>The default is 10.</p>
CISPLITR= CI	<p>Defines the default minimum ratio of Control Interval splits to every 100 Control Interval's in a VSAM KSDS.</p> <p>The default is 10.</p>
CKPTALLOC= CKPTA	<p>Defines the allocation method to use when allocating the checkpoint data set.</p> <p>SMS – Allocate by SMS storage class.</p> <p>UNIT – Allocate by unit name.</p> <p>Note – If SMS is specified, a valid SMS storage class name must also be specified for CKPTSTORCLASS. If UNIT is specified, a valid unit name must also be specified for CKPTUNIT.</p> <p>The default is UNIT.</p>
CKPTDATACLASS= CKPTD	<p>Specifies the SMS data class to use when CKPTALLOC is SMS.</p>
CKPTMGMTCLASS= CKPTM	<p>Specifies the SMS management class to use when CKPTALLOC is SMS.</p>
CKPTPREFIX= CKPTP	<p>Specifies the high level qualifier to use when constructing the name of the checkpoint file. To make it possible for the RECOVER command to find the checkpoint file for a failed REORG, the checkpoint data set is cataloged with a name of</p> <p style="text-align: center;">&CKPTPREFIX.REORGCKP.jobname.Dccyyddd.Thhmms</p> <p>You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.</p> <p>The default is FDRREORG.</p>
CKPTSTORCLASS= CKPTS	<p>Specifies the default SMS storage class to use when CKPTALLOC is SMS.</p>
CKPTUNIT= CKPTU	<p>Defines a valid unit name containing disk devices to use when CKPTALLOC is UNIT.</p> <p>The default is SYSDA.</p>

31.05 CONTINUED . . .

ENQERR=

Specifies if FDRREORG should set a return code of 8, 4 or 0, if a data set selected for reorganization is in use by another job or user.

YES – Set return code 8 if no data sets are selected.

WRN – Set return code 4 if no data sets are selected.

NO –Set return code 0 if no data sets are selected.

**IAMCOMPPERM=
IAMCOMPP**

Specifies if FDRREORG should backup compressed IAM data sets in compressed format when BACKUP=GDG or BACKUP=PERM has been specified.

YES – Backups will be in compressed format.

NO – Backups will not be in compressed format.

The default is NO.

Note – This feature requires IAM 6.3 or above and is only supported for files in enhanced format.

**IAMCOMPTEMP=
IAMCOMPT**

Specifies if FDRREORG should backup compressed IAM data sets in compressed format when BACKUP=TEMP has been specified.

YES – Backups will be in compressed format.

NO – Backups will not be in compressed format.

The default is YES.

Note – This feature requires IAM 6.3 or above and is only supported for files in enhanced format.

**IAMDEFINE=
IAMDEF**

Specifies the default option for deleting and defining IAM data sets that will be reorganized.

YES – IAM data sets will be deleted and defined before reloading.

NO – IAM data sets will not be deleted and defined before reloading.

The default is NO.

Note – This feature requires IAM 6.3 or above. Additionally, files defined with an earlier release of IAM will not be redefined by FDRREORG until they are first redefined with IAM 6.3 or above.

KEEPCRDATE=

Specifies if FDRREORG should retain the original creation date when redefining VSAM.

YES – FDRREORG will retain the original creation date.

NO – FDRREORG will use the current date.

The default is YES.

31.05 CONTINUED . . .

**FREESPACE=
FREESP**

Defines the default selection cut off for VSAM CA or CI free space percent. This value is used to prevent reorganizations of VSAM KSDS's or AIX's that have very high free space percentages. To reduce splits for some files in online systems, some VSAM files are loaded with little or no free space and then altered to have a very high CA and/or CI freespace percent. If these files are processed by FDRREORG, the current freespace percentage is used resulting in a dramatic increase in space.

The default is 50.

**LASTAPEPREFIX=
LASTAPEP**

Specifies the high level qualifier to be used when cataloging or locating the special LASTAPE catalog entries. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

The default is FDRREORG.

**LISTNOREORG=
LISTNOR**

Specifies the default option for printing the NOREORG list for REORG or SIMULATE functions.

YES – The NOREORG list is printed for REORG or SIMULATE functions.

NO – The NOREORG list is not printed.

The default is YES.

**LOGALLOC=
LOGA**

Defines the allocation method to use when allocating the log data set.

SMS – Allocate by SMS storage class.

UNIT – Allocate by unit name.

Note – If SMS is specified, a valid SMS storage class name must also be specified for LOGSTORCLASS. If UNIT is specified, a valid unit name must also be specified for LOGUNIT.

The default is UNIT.

**LOGDATACLASS=
LOGMGMTCLASS=**

Specifies the SMS data class to use when LOGALLOC is SMS.

Specifies the SMS management class to use when LOGALLOC is SMS.

**LOGPREFIX=
LOGP**

Specifies the high level qualifier to use when constructing the name of the log file. To make it possible for the RECOVER command to find the log file, the log data set is cataloged with a name of

&LOGPREFIX.REORGLOG.jobname.Dccyyddd.Thhmmss

You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF user id or RACF group id be used.

The default is FDRREORG.

LOGSTORCLASS=

Specifies the SMS storage class to use when LOCALLOC is SMS.

31.05 CONTINUED . . .

LOGUNIT=	<p>Defines a valid unit name containing disk devices to use when LOGALLOC is UNIT.</p> <p>The default is SYSDA.</p>
MAXENQ=	<p>Specifies the maximum number of outstanding enqueue's allowed for the DSNRETRY ENQ and WAIT options for each volume processor task. When this limit is reached, data sets will be added to task retry queue as if the RETRY option was specified. The enqueue will be issued when the outstanding enqueue count falls below this maximum. Once the volume processor task has completed processing all volumes and is waiting for data sets queued with the WAIT option, no additional enqueue's will be issued.</p> <p>The default is 999.</p>
MAXENVERR= MAXENV	<p>Specifies the maximum number of environmental errors allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set. No additional REORG or SIMULATE commands will be processed. Environmental errors are any backup or re-load failures not caused by a system abend. Insufficient space, or target data set not available, are examples of environmental errors.</p> <p>The default is 99.</p>
MAXFILE=	<p>Specifies the maximum number of backup files to place on a single tape volume.</p> <p>The default is 255.</p>
MAXSYSERR= MAXS	<p>Specifies the maximum number of system abends allowed. When this limit is reached, all subtasks will terminate after processing the currently active data set. No additional REORG or SIMULATE commands will be processed.</p> <p>The default is 99.</p>
MAXTASKS= MAXT	<p>Specifies the maximum number of concurrent volumes to be processed. You may specify any number from 1 to 15, inclusive. Please note that the actual number of subtasks possible is limited by the amount of available virtual storage. For SIMULATE commands with MAXTASKS=15, a region size of 2.5 megabytes should be sufficient on an XA or ESA system. For non-XA systems, a region size of 5 megabytes should be sufficient. For REORG commands with MAXTASKS=15, a region size of 8.5 megabytes should be sufficient on an XA or ESA system. For non-XA systems, MAXTASKS=15 is not possible because the region size required exceeds 16 megabytes. For non-XA systems, the default value of 4 is probably the highest value that can be specified if the maximum region size available is used.</p> <p>The default is 4.</p>

31.05 CONTINUED . . .

MSGLEVEL=
MSGL

Specifies the lowest level message type to be displayed on the subtask print file which has the ddname REORGPRT.

I – Informational, warning, and error messages will be displayed.

W – Warning and error messages will be displayed

E – Only error messages will be displayed.

The default is I.

MSGTIMESTAMP=
MSGT

Specifies whether messages written to REORGPRT will be suffixed with the current date, time, and the internal subtask id that issued the message. If YES is requested or defaulted, the LRECL of the REORGPRT file is increased to 151.

YES – Messages will be timestamped.

NO – Messages will not be timestamped.

The default is YES.

NOUPDATES=
NOUP

Specifies the default action to take for IAM or VSAM data sets selected for reorganization that have had no adds, no deletes and no updates.

YES – The data set will be reorganized.

NO – The data set will not be reorganized.

The default is NO.

OFULL=
OF

Defines the default minimum percentage of independent overflow used in an IAM file.

The default is 80.

OVERFLOWINDEX=

Specifies the default minimum amount of memory in bytes required for the in storage IAM overflow index.

The default is 1048576 (1 megabyte).

OWNERSTRING=
OWN

Specifies a string of up to four characters which will be used as an eye catcher in the ownerid of the data component of a VSAM KSDS. REORG will update the ownerid field of the data component's catalog entry to record information used to insure data set integrity and to identify data sets that were being processed by REORG when either a system failure occurred, or REORG was canceled or otherwise terminated. See the RECOVER command for additional information.

The default is FDR\$.

PCTTRECO=

Defines the default minimum percentage of total records in an IAM file that are in the IAM independent overflow area.

The default is 10.

PDSFULL=
PDSF

Defines the default minimum percentage of allocated space used by a PDS.

The default is 90.

PDSEXTENTS=
PDSEX

Defines the default minimum number of extents for a PDS.

The default is 17.

31.05 CONTINUED . . .

PEFULL= PEF	<p>Defines the default minimum percentage of prime extension used in a IAM file.</p> <p>The default is 100.</p>
PEUDATAR= PEUD	<p>Defines the default minimum ratio times 100 of used prime extension blocks to prime data blocks in an IAM file.</p> <p>The default is 100.</p>
POOLDASD=	<p>Specifies if POOLDASD from Boole & Babbage or a similar allocation control product is installed, and unit names which have not been defined to the system are being used by POOLDASD to control volume allocation.</p> <p>YES – Accept undefined unit names.</p> <p>NO – Do not accept undefined unit names.</p> <p>The default is NO.</p>
RECORDCOUNTS=	<p>Specifies if FDRREORG should issue a message containing the number of records read and the number of records loaded after reorganizing IAM and VSAM data sets.</p> <p>YES – Issue FDRREORG message FDRS55 which contains the record counts.</p> <p>NO – Do not issue the FDRS55 message.</p> <p>The default is NO.</p>
REDUCEPRIMARY=	<p>Specifies if FDRREORG should decrease the primary space allocation of enhanced format IAM data sets or VSAM KSDS's if the total allocated space is less than the primary space allocation used to originally define the file. This option only applies to data sets which are redefined by FDRREORG.</p> <p>Specifying YES for this option will enable FDRREORG to preserve the current space allocation for data sets which may have had excess free space released by COMPAKTOR, or for IAM files loaded with RELEASE=YES.</p> <p>YES – FDRREORG will use the total allocated space as the primary when defining data sets that have had the total allocated space reduced to an amount less than the original primary.</p> <p>NO – FDRREORG will always use the original primary space allocation.</p> <p>The default is YES.</p> <p>Note: FDRREORG will also reduce the size of the index of a VSAM KSDS if was defined with NOIMBED.</p>

31.05 CONTINUED . . .

SELTERR=

Specifies if FDRREORG should set a return code of 8, 4 or 0 if no data sets are selected by REORG or SIMULATE. Due to the nature of the selection process and the selection criteria specified, this may be a natural occurrence.

YES – Set return code 8 if no data sets are selected.

WRN – Set return code 4 if no data sets are selected.

NO – Set return code 0 if no data sets are selected.

The default is YES.

SORTRPT=

Specifies the default sorting option for the report written to REORGRPT. Specifying YES requires that information related to each data set processed or selected be kept in storage until the function is completed. Because this might require a significant amount of virtual storage, this option is only supported for MVS/XA, MVS/ESA or OS/390 systems.

YES – Report is produced at the end of a REORG, SIMULATE or RECOVER function in data set name and volser order.

NO – Report is produced as data sets are processed or selected.

The default is NO.

**SUBTASKABEND=
SUBTASKA**

Specifies the action to take if a volume processor subtask abends.

CONT – Continue processing without the subtask.

TERM – Quiesce all active work and terminate.

The default is CONT.

UPDATEDPDS=

Specifies the default action to take for partitioned data sets that do not have a current backup. If the update indicator is on in a data sets format 1 DSCB, the data set is considered to not have a current backup.

YES – The data set will be compressed.

NO – The data set will not be compressed.

The default is YES.

**VSAMDEFINE=
VSAMDEF**

Specifies the default option for deleting and defining VSAM KSDS's that will be reorganized.

ALWAYS – VSAM KSDS's are always deleted and defined before reloading.

IFREQ – Only VSAM KSDS's that cannot be reused will be deleted and defined before reloading.

NO – Disables delete and define of all VSAM KSDS's. KSDS's that cannot be reused will not be reorganized.

The default is IFREQ.

31.06 FDRREORG JCL EXAMPLES

DISPLAY Display the current FDRREORG options.

OPTIONS

```
//PRINT      EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
PRINT
```

DISPLAY Display the current FDRREORG NOREORG list.

NOREORG

```
//PRINT      EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
PRINT NOREORG
```

CHANGE Change the checkpoint and log file high level qualifier to the RACF user id of the submitter.

PREFIXES

```
//ZAP        EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
ZAP CKPTPREFIX=&RACFUID,
LOGPREFIX=&RACFUID
```

CHANGE Disable the default selection criteria for all data set types.

OPTIONS

```
//ZAP        EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
ZAP CASPLITR=0,CISPLITR=0,FREESPACE=100,
OFULL=0,PCTTRECO=0,PEFULL=100,PEUDATAR=100,
PDSFULL=0,PDSEXTENTS=1
```

UPDATE Add SYS2 data sets to the NOREORG list.

NOREORG BY DATASET

```
//NOREORG    EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
NOREORG DSN=SYS2.**
```

UPDATE Add all IPLxxx and CATxxx volumes to the NOREORG list.

NOREORG BY VOLUME

```
//NOREORG    EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
NOREORG VOL=(IPL*,CAT*)
```

UPDATE Add SYS3 data sets on volume SYSVOL to the NOREORG list.

NOREORG

```
//NOREORG    EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
NOREORG DSN=SYS3.** ,VOL=SYSVOL
```

CONTINUED ...

31.06 CONTINUED . . .

RESET Reset the FDRREORG options to the distribution settings.
OPTIONS

```
//RESET      EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
RESET
```

RESET Remove all user modifiable entries from the NOREORG list.
NOREORG

```
//RESET      EXEC  PGM=FDRREOZO
//SYSPRINT   DD    SYSOUT=*
//SYSLIB     DD    DSN=fdrreorg.library,DISP=SHR
//SYSIN      DD    *
RESET NOREORG
```


THIS PAGE WAS INTENTIONALLY LEFT BLANK

32.01 FDRCOPY REORG STATEMENT

```

REORG          TYPE=DSF
SIMREORG

,DSNENQ=NONE|USE|TEST|HAVE

,ENQERR=NO

,ENQERR=BYPASS|PROCESS

,MAXCARDS=nnnnn

,MAXERR=nnnn

,SELTERR=NO|YES

```

FDRCOPY REORG PDS reorganization can also be invoked directly under program FDRCOPY. See the examples in [Section 32.03](#) for the JCL required.

REORG STATEMENT The REORG statement invokes FDRCOPY PDS reorganization (compression in-place), similar to the PDS compression of IEBCOPY, but with 50% to 90% less elapsed time and similar reduction in other resources used. For all DSORG=PO (partitioned) data sets selected, REORG will relocate all active members toward the beginning of the data set, leaving the maximum amount of unused space at the end of the PDS for new and updated members. An unlimited number of PDSs may be reorganized in execution of FDRCOPY from any number of disk volumes.

SIMREORG STATEMENT The SIMREORG statement simulates the operation of the REORG statement. SIMREORG must still read the input data sets to determine member location, but WRITES are simulated and the data sets will not be modified. SIMREORG will report on all selected data sets, indicating the number of tracks that a real REORG will reclaim.

LICENSING REQUIREMENTS Although REORG is part of FDRCOPY, the REORG function is enabled only if your installation is licensed for FDRREORG, a separately-priced component of FDR.

FDRREORG ([see Section 30](#)) reorganizes VSAM data sets, IAM (Innovation Access Method) data sets, and PDSs (using FDRCOPY for the PDS reorganization). If you execute PDS reorganization through FDRREORG, it enhances the reorganization process by providing additional selection criteria, retrying unavailable data sets, etc.. However, these enhancements may add additional overhead; executing FDRCOPY REORG directly will provide the fastest PDS reorganization.

CONTINUED . . .

32.01 CONTINUED . . .

OPERANDS	TYPE=DSF	Required. Specifies that this is a data set operation.
	DSNENQ=	<p>Specifies whether all of the data sets selected for reorganization will be ENQed. See "Data Set Enqueue Option" in Section 21.02 for more details.</p> <p>If the ENQ fails on an data set, meaning that some other task has the data set enqueued, a FDR158 warning message is issued for the data set and it is bypassed unless the ENQERR=PROCESS operand is specified. A successful ENQ will prevent any other task from using the data set until the reorganization is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be reorganized. The options for DSNENQ= are:</p> <p>USE -- The data sets will be enqueued for the duration of the reorganization from this disk volume. This is the most frequently used option and is the default for FDRCOPY.</p> <p>TEST -- The data sets will only be tested to see if they are enqueued to another task at the time that the reorganization on this volume starts.</p> <p>HAVE -- The data sets will be enqueued for the duration of the reorganization on this disk volume. If not available, a message (FDRW27) is issued to the MVS operator. See "Data Set Enqueue Option" in section 21.02 for the valid responses.</p> <p>NONE -- No data set ENQ will be issued.</p> <p>CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.</p> <p>Default: USE.</p> <p>Recommendation: use DSNENQ=USE (the default) or HAVE if you want to be sure that no other task uses the data set until the reorganization is complete. However, use DSNENQ=NONE when another data set by the same name on another volume may be in use (e.g., a duplicate name on an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.</p>
	ENQERR=	<p>NO – If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see "Step Termination" in Section 21.02). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally.</p> <p>Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.</p>

32.01 CONTINUED . . .

- ENQERR=** Specifies processing if the DSNENQ= option finds that an input data set is in use (enqueued):
- BYPASS** – do not reorganize an active data set.
- PROCESS** – reorganize a data set even if it is active (a warning message will still be produced).
- Default: BYPASS.
- Note- both ENQERR=NO and ENQERR=BYPASS/PROCESS may be specified on the same statement.
- MAXCARDS=** Accept additional SELECT/EXCLUDE statements (over 250).
- Default: 250 statements.
- MAXERR=** Specifies the number of read or write disk errors that are to be bypassed by FDRCOPY prior to ABENDING the operation. MAXERR may specify a value from 1 to 9999 errors. Read and write errors are counted separately, but MAXERR= specifies the maximum for each counter.
- Default is 20 errors.
- SELTERR=** Specifies what will happen at step termination if FDRCOPY finds that a SELECT or EXCLUDE statement was never referenced (no data set on any disk was selected by the statement):
- NO** – a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.
- YES** -- a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.
- Default: YES unless overridden in the FDR Global Option Table ([See Section 90](#)).

32.02 FDRCOPY SELECT/EXCLUDE STATEMENT FOR REORG

SELECT	DSN=filter	,EMPTY=YES NO
S	CATDSN=filter	
	DD=ddname	
	ALLDSN	,LIST=YES NO
EXCLUDE	,CATALOG=catname	
X	,MCATALOG=catname	,PRTALIAS
	,DSNENQ=NONE	,VOL=vvvvvv
	,DIRBLKS=nn	
	,%DIRFREE=nn	

**SELECT/
EXCLUDE
STATEMENT
for REORG**

The SELECT statement selects the PDS (partitioned) data sets that will be reorganized. FDRCOPY will scan the VTOC of each input volume for data sets which match the parameters on a SELECT statements; data sets which are not DSORG=PO (PDS) will be bypassed. Data sets may be selected by fully-qualified data set name or by using generic data set name selection (DSN=), or by copying a data set name from a DD statement (DD=). All data sets on an input volume may be selected (ALLDSN). If the SELECT statement specifies the VOL= operand, the named volumes are automatically processed as input volumes.

Data sets may also be selected from the system catalogs by fully-qualified name or using a generic filter (CATDSN=). The volumes found in the catalog entries will be automatically processed as input volumes, and only the data set name found in the catalogs will be selected from their VTOCs.

The EXCLUDE statement prevents certain data sets from being reorganized. The data sets to be excluded may be specified by fully-qualified name or by using generic data set name selection (DSN=) or copied from a DD statement (DD=). The EXCLUDE statement may be used to exclude particular data sets that would be selected by a more-encompassing SELECT statement. Since SELECT/EXCLUDE statements are scanned in the order they are input, EXCLUDE statements should usually precede SELECT statements.

SELECT and EXCLUDE statements will apply to all input disk volumes unless the VOL= parameter is specified. OPERANDS

OPERANDS

DSN= Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the VTOCs of selected input volumes.

EXAMPLES : DSN=USER1.JCL.CNTL
 DSN=**LIST
 DSN=PROD++.**.LIB*

DD= Specifies that a data set name is to be taken from a DD statement. This operand must point to the DDNAME of a JCL statement. Using this option enables the user to specify a non-standard data set name.

EXAMPLE : SELECT DD=DD1
 //DD1 DD DSN=USER.JCL,DISP=SHR

CONTINUED . . .

32.02 CONTINUED . . .

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= parameter (if there are no DISKx DD statements pointing to those volumes, FDRCOPY will dynamically allocate and process them as input volumes).

If a filter is specified, then ICF catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT CATDSN=dsname was present for each of them. It may be necessary to specify MAXCARDS=nnnnn if a large number of data sets are selected by the filter. Data sets which are obviously not PDSs (such as VSAM and GDGs) will not be selected from the catalog, but other non-PDS entries can not be excluded until later when the VTOC is processed on each volume. Additional considerations for CATDSN=filter are explained in [Section 80](#).

CATDSN= is supported only on SELECT statements. However, a preceding EXCLUDE statement with DSN= and/or VOL= can exclude data sets from selection by CATDSN=.

If the VOL= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

```
EXAMPLES :      CATDSN=USER1.JCL.CNTL
                  CATDSN=**COBOL
                  CATDSN=PROD++.**.LIB*
```

Normally CATDSN= will not display the data sets it selects from the catalogs, you will see the names only when FDRCOPY actually finds and selects the data sets in the VTOCs of the volumes they are cataloged to. To display all of the data sets selected specify PCATDSN=filter. However, PCATDSN may display data sets which are not selected because they are not PDSs.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

ALLDSN

Specifies that FDRCOPY is to reorganize all PDSs on the volumes specified. DSN=** is equivalent to ALLDSN.

NOTE: DSN=, CATDSN=, DD= and ALLDSN are mutually exclusive. One and only one of these operands must be specified on each SELECT or EXCLUDE card.

CATALOG=
MCATALOG=

Specifies the name of a user catalog (CATALOG=) or alternate master catalog (MCATALOG=) to search when CATDSN= is specified. [See Section 80](#) for details.

Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter.

32.02 CONTINUED . . .

DSNENQ= **NONE** -- the data set enqueue will not be done for the selected data sets. This can be used to override the DSNENQ= operand on the REORG statement for certain data sets that you know will probably be ENQed by another task. Default: enqueue option is determined by the DSNENQ option specified on the COPY or MOVE statement.

WARNING: DSNENQ=NONE will allow PDSs to be reorganized even while they are in use, which may cause failures in jobs or users who are reading them, and may cause FDRCOPY to fail if they are being updated.

DIRBLKS= **nn** -- specifies that all PDSs processed by this SELECT statement will have their directories expanded by 'nn' directory blocks during the reorganization. 'nn' may be from 1 to 99. DIRBLKS= is designed to be used for a one-time expansion, if left on the SELECT it will expand the directories every time the REORG is run.

%DIRFREE= **nn** -- specifies that any PDS processed by this SELECT will have its directory expanded if less than 'nn' percent of the directory is currently unused (based on the number of directory blocks in use and the bytes used in the last or only block). Enough directory blocks will be added so that it does have nn% free directory space after REORG. For example, on a PDS with 95 directory blocks (5 blocks free), %DIRFREE=10 will expand it to 100 blocks (10 blocks free). 'nn' may be from 1 to 99. %DIRFREE= may be permanently specified on the SELECT since it will only expand directories which are short on free space.

The default, if neither DIRBLKS= nor %DIRFREE= is specified, is that the directory size of selected PDSs will be unchanged.

For both DIRBLKS= and %DIRFREE=, a maximum of one track's worth of directory blocks (e.g., 45 on a 3390) will be added to any PDS in one REORG execution; it is possible that several REORG executions will be necessary to fully expand some directories. Also, the directory will never be expanded beyond the end of the first extent (an MVS restriction), and the number of blocks added may be reduced (possibly to zero) if little or no unused space exists in the PDS before reorganization (the PDS itself will never be extended to make room for the expanded directory).

EMPTY= **YES** -- empties the directory of the selected PDSs, so that they appear to have no members. The Last Block Pointer in the DSCB will point immediately after the directory. This may be used to empty PDSs without having to delete all members and reorganize (or delete and reallocate the PDS). This can be useful for emptying linklist data sets and others which are always active (DSNENQ=NONE may be required to REORG active data sets).

NO -- selected PDSs will be reorganized normally.

The default is NO.

WARNING: EMPTY=YES should be used with caution since all members in the selected PDSs will be DELETED!

32.02 CONTINUED . . .

- LIST=** **YES** -- For each selected PDS, all member names in the PDS are listed, indicating which members were moved.
- NO** -- Member names are not listed.
- Default is NO.
- PRTALIAS** When used on a SELECT statement with CATDSN= will display all of the alias names and user catalogs that were searched. This is effective only when the data set name mask you provided forces CATDSN= to start in the master catalog and search one or more alias catalogs. If the non-mask characters at the beginning of the mask are sufficient to cause CATDSN= to begin its search in a user catalog, PRTALIAS is ignored.
- VOL=** Specifies the input disk volume serial(s) to which this SELECT/EXCLUDE statement is to apply. It may specify a single serial (e.g., VOL=ABC123) or a group of volumes all starting with the same prefix (e.g., VOL=ABC*). If there are online disk volumes matching the VOL= value that are not pointed to by DISKx DD statements in the FDRCOPY step, FDRCOPY will automatically dynamically allocate them and process them. If specified on a SELECT CATDSN=, only PDSs cataloged to the specified volumes are chosen.
- Default: the SELECT/EXCLUDE statement will apply to all input volumes and SELECT CATDSN= will select data sets regardless of the cataloged volume serial.

32.03 FDRCOPY REORG EXAMPLES

Reminder: the PDS reorganization functions shown here will not work unless your installation is licensed for FDRREORG, an extra-cost feature of FDR .

REORGANIZE ONE PDS A single PDS will be reorganized. Its volume will be extracted from the catalog. All member names will be listed showing which were moved.

```
//REORG      EXEC   PGM=FDRCOPY , REGION=2M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
      REORG      TYPE=DSF
      SELECT     CATDSN=USER1.JCL.CNTL , LIST=YES
/*
```

REORGANIZE ONE PDS This is a cataloged procedure that can be used to reorganize one or more PDS data sets. The PDS name or a filter name is a symbolic parameter to the proc. The FDRCOPY control statements are passed in the EXEC statement PARM, so SYSIN is a DUMMY data set.

```
//REORG      PROC   PDS='NOT.GIVEN'
//REORG      EXEC   PGM=FDRCOPY , REGION=2M ,
//            PARM=' REORG TYPE=DSF /SELECT ' CATDSN=&PDS '
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     DUMMY
//            PEND
```

Example: To reorganize all data sets with a high level index of USER1 and a last index level of LOAD specify:

```
//REORG      EXEC   REORG , PDS='USER1.**.LOAD '
```

To duplicate the first example with this proc, use:

```
//REORG      EXEC   REORG , PDS='USER1.JCL.CNTL '
```

REORGANIZE PDSs ON SPECIFIED VOLUMES All PDSs on volumes starting with TSO will be reorganized. Data sets currently allocated to another job or user will be bypassed.

```
//REORG      EXEC   PGM=FDRCOPY , REGION=4M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
      REORG      TYPE=DSF
      SELECT     DSN=**, VOL=TSO*
/*
```

REORGANIZE MANY PDSs A number of cataloged PDSs, which may be on many volumes, will be reorganized. MAXCARDS= is specified in case over 250 data sets are selected. Data sets currently allocated to another job or user will be bypassed. Member names will be listed for all JCL libraries.

```
//REORG      EXEC   PGM=FDRCOPY , REGION=4M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
      REORG      TYPE=DSF , MAXCARDS=1000
      SELECT     CATDSN=**.CNTL
      SELECT     CATDSN=**.JCL , LIST=YES
      SELECT     CATDSN=PROD.**
/*
```

CONTINUED . . .

32.03 CONTINUED . . .

REORGANIZE ACTIVE PDSs REORG will normally bypass data sets that are ENQed (in use) by another job or user, but that ENQ is by data set name, not volume serial. This job will reorganize PDSs on an alternate SYSRES volume even though a data set by the same name is in use on the active SYSRES.

WARNING: ENQERR=PROCESS (or DSNENQ=NONE) must not be used to reorganize PDSs which are actually in use; they may become unusable.

```
//REORG      EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
              REORG   TYPE=DSF,ENQERR=PROCESS
              SELECT   DSN=SYS**,VOL=ALTRES
/*
```

SIMULATE REORGANIZATION A simulated reorganization of all PDSs on all online volumes will be done. Since FDRCOPY cannot determine the results of the reorganization without knowing the contents of the PDS, each selected PDS will be read, but all writes will be bypassed and no PDS will be changed. Note: This example may run for a long time since FDR will read all the PDS's in your installation.

```
//REORG      EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
              SIMREORG TYPE=DSF
              SELECT   DSN=**,VOL=*
/*
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

33.01 FDRREORG ISPF INTERFACE

For ABR customers, FDRREORG is supported by the ABR SRS dialogs on ISPF. As described in [Section 54.50](#), SRS allows end users to generate a table of selected data sets and their characteristics, and to enter commands to be executed against those data sets.

**FDRREORG
PANEL**

If the "FDRREORG" command is entered for a data set, you will receive a panel similar to this:

```

----- FDRREORG -----
COMMAND ==>
  Edit generated JCL      Submit generated JCL      SCROLL ==> HALF
Operands for REORG/SIMULATE statement (section 30.03):
==> REORG NODEFAULTS
-----
DSNAME / Filter ==> 'BAB.JCL.CNTL'                  also ALLDSN
Volume Serial   ==> TSOWK2                          optional
Storage Group   ==>                                optional
Operands for SELECT statement (section 30.05)
==>

```

The name and volser of the input data set will be filled in. Since you chose the data set to be reorganized, the NODEFAULTS operand is assumed so that FDRREORG will not test its characteristics and possibly bypass it. This can be used on any data set supported by FDRREORG ([See Section 30.01](#)).

REORG PANEL

If the "REORG" command is entered for a data set, you will receive a panel similar to this:

```

----- FDRSRS - Reorg Data Set ----- Row 1 to 1 of 1
COMMAND ==>
  Edit generated JCL      Submit generated JCL      FG - execute in the foreground
Operands for REORG TYPE=DSF statement (section 32.01):
==> REORG TYPE=DSF
DSNAME / Filter ==> 'BAB.JCL.CNTL'
Volume Serial   ==> TSOWK2
Add dir. blocks ==>
Operands for SELECT DSN= statement (section 32.02):
==> NOTIFY=BAB

```

The name and volser of the input data set will be filled in. This invokes FDRCOPY to reorganize a PDS, so it will not work on other data set types.

**PANEL
EXECUTION**

If needed, you can provide additional operands for the REORG statement or the SELECT statement on the lines indicated.

For FDRCOPY REORG, you can submit a batch FDRCOPY jobstream to perform the copy/move (EDIT or SUBMIT commands) or to attach FDRCOPY under TSO and execute the copy/move in the foreground (FG). The foreground operation will tie up your TSO session until the operation is done.

For FDRREORG, you must submit a batch jobstream.

**ISPF
INSTALLATION
PANELS**

The ISPF panels used for setting FDRREORG options and maintaining the FDRREORG option table are documented in [Sections 90.23](#) and [90.31](#).

THIS PAGE WAS INTENTIONALLY LEFT BLANK

40.01 COMPAKTOR INTRODUCTION

CPK COMPAKTOR's primary function is to reduce free space fragmentation on a volume. In addition, CPK can combine non-VSAM multi-extent data sets and VSAM multi-extent components into one extent. CPK provides the capability of enlarging and/or moving the VTOC.

CPK can release unused space within PS, PO and PDSE data sets and VSAM components, or some of that free space can be retained for expansion of the data set. CPK can free excess space as part of a volume reorganization or as an independent function.

The FASTCPK option of COMPAKTOR consolidates free space, merges extents, positions data sets and releases space from volumes **without taking an FDR backup**.

FASTCPK achieves the almost same level of free space consolidation and reduction of multi-extent data sets as regular COMPAKTion that requires a backup. FASTCPK accomplishes this without compromising the integrity of your data. FASTCPK maintains a log of its processing and provides an automatic recovery function which allows COMPAKTion to be restarted if the operating system or the job fails.

FASTCPK performs the reorganization in **50 to 90%** less elapsed and CPU time than a regular COMPAKTion due to in-place COMPAKTion, its algorithms and I/O performance options. FASTCPK can COMPAKT many volumes in a single execution with simple control statements, e.g., CPK TYPE=FASTCPK,VOL=(TSO*,WORK*)

FASTCPK can be run daily with conditional statements that will only select volumes requiring COMPAKTion (those exceeding a number of free space areas or IBM's Fragmentation Index).

FASTCPK WILL COMPAKT VOLUMES TO A SMALLER NUMBER OF FREE SPACE AREAS IN A FRACTION OF THE TIME REQUIRED BY COMPETING PRODUCTS.

CPK provides both detail and summary reports on the condition of the volume, space usage and fragmentation, and data set types and placement. Summaries from before and after COMPAKTion are displayed side-by-side. COMPAKTOR also has a simulation capability to provide a preview of the improvements that it can make on the volume, or to test the result of various options. A combined summary report displays a few lines per volume processed, allowing you to quickly see the result of the COMPAKTion and its estimated time.

CPK can reorganize an active volume (with data sets currently in use). CPK will determine which data sets are currently active, bypassing moving or releasing space on these data sets. All inactive data sets will be enqueued on, preventing their use during the COMPAKTion or release function.

COMPAKTOR can be licensed as a separate product without a requirement that you be licensed for FDR.

**NEW FASTCPK
USERS AND
FASTCPK
TRIALS**

Go to [Section 40.23](#) and use Example 5 as a model on how to use FASTCPK. A FASTCPK simulation only takes a few seconds. After you have run a few simulations, contact Innovation Technical Support to go over the results.



FASTCPK significantly changes the way that COMPAKTOR users should use the product. All COMPAKTOR jobs should be converted to use FASTCPK (except when moving or enlarging the VTOC). See [Section 40.25](#) "CPK Hints and Recommendations" for suggestions on running FASTCPK, and answers for common questions and problems.

FASTCPK anticipated results and an accurate estimate of the elapsed time required in a real COMPAKTion can be obtained by running a FASTCPK simulation. See [Section 40.23](#) for FASTCPK simulation examples.

CONTINUED . . .

40.02 COMPAKTOR BACKGROUND

The continuing increase in the number and size of direct-access devices in most installations has created a number of problems and requirements. Many can be resolved by periodically reorganizing disk volumes.

FREE SPACE FRAGMENTATION Many a disk volume has small non-contiguous areas of free space spread haphazardly throughout the volume. Although the primary allocation of a new data set can be satisfied in multiple extents if no single free space extent is large enough, this may force the data set to be allocated from two or more non-contiguous free areas, affecting its performance.

MULTIPLE EXTENTS Either through natural growth or inaccurate space estimates, data sets tend to grow beyond their original allocated data space. They may end up owning anywhere from 2 to 16 different extents (up to 123 extents for ICF VSAM components and some SMS-managed data sets). It may become impossible to extend the data set even though a large amount of free space is available on the volume.

WASTED SPACE Often, many non-volatile data sets occupy a larger data space than necessary; this is caused by inaccurate data space requirement estimates. The result is to effect a net loss in available free space on the involved volumes.

SMALL VTOCs Sometimes, a VTOC becomes full and it becomes impossible to allocate data sets on that volume, even though free space is available.

INACCURATE VTOCs Hardware/software system failures may lead to inaccurate VTOCs which do not reflect the true status of their volumes. This is a dangerous condition and can cause loss of data.

SYSTEMS PERFORMANCE The above mentioned symptoms, when prevalent, tend to deteriorate a system's performance; total system throughput is degraded and TSO and online systems evidence poor response times. Moreover, this deterioration often occurs imperceptibly, over some period of time; thus, you may have difficulty in determining its causes. In some cases, it can lead to premature hardware upgrades.

WHY DISK VOLUMES ARE SELDOM REORGANIZED Although most sophisticated users recognized and understood these problems and the need to reorganize disk volumes, in practice they seldom did so. The reason is that, without COMPAKTOR, reorganizing a disk volume is a slow, laborious and error-prone process; very few users thought that the benefits of reorganization outweighed its difficulties.

CONVERSION BETWEEN DISK MODELS Within a given type of disk, such as 3390, several models are available with varying capacities, and installations often need to convert volumes from one model to another. However, because of the different capacities, the ideal locations for the VTOC and the size of the VTOC vary from one model to another. During the conversion, users usually want to move and expand (or shrink) the VTOC, and to position VTOC indexes and VVDSs near the VTOC. (Note that CPK cannot be used to convert between different types of disks, such as 3380 to 3390; [see sections 20 and 21](#) to use FDRDSF and FDRCOPY for unlike device conversion).

The following sections of this manual describe how INNOVATION DATA PROCESSING has solved these problems, with its COMPAKTOR program product.

CONTINUED . . .

40.03 COMPAKTOR OVERVIEW

COMPAKTOR provides an efficient and easy-to-use method of reorganizing a disk volume, positioning data sets or releasing unused space within data sets. It does this via a simple one-step process, with a minimum of JCL and control statements required.

FEATURES **REDUCTION OF FREE SPACE FRAGMENTATION.** On most volumes, CPK drastically reduces the number of free space areas, through contiguous placement of data sets.

MERGING OF DATA SET EXTENTS. Most multi-extent data sets will be merged into a single extent by COMPAKTOR.

DATA SET POSITIONING. COMPAKTOR has a flexible and convenient feature for user-specified data set placement.

ELIMINATION OF UNUSED DATA SET SPACE. A user may free all or part of the unused space within physical sequential (PS), partitioned (PO), and/or ICF VSAM data sets. The freeing of space can be done during COMPAKTion or as a independant option of CPK (TYPE=RLSE).

VTOC POSITION AND SIZE. COMPAKTOR allows you to change the position and size of the VTOC. This can only be done by a COMPAKT-from-backup, not a FASTCPK.

VTOC ANALYSIS. COMPAKTOR always performs a complete analysis of a VTOC, often detecting structural VTOC errors not detected by any IBM software. Any errors detected result in diagnostic messages. Under some conditions, CPK even corrects VTOC errors.

ERROR FREE VTOCS. After a successful COMPAKTion, all reorganized disk volumes have error-free VTOCs. The volume free space (in Format 5 DSCBs or the indexed VTOC) will be accurate.

VOLUME MAPS. CPK provides track maps of disk volumes, in track sequence. Maps of a volume both before and after COMPAKTion are automatically provided. CPK also provides a summary report on the status of the volume.

MAPPING FEATURE. As an option, you can request that CPK only map a volume. The volume may be a disk volume residing on a direct-access device or an FDR backup of a disk volume residing on a tape or sequential file on disk. Full VTOC analysis is still performed.

SIMULATION FEATURE. This feature provides full maps of a disk volume both before and after reorganization or release. However, no actual COMPAKTion is performed and the volume is never modified.

MULTIPLE VOLUME FUNCTIONS. For most functions, COMPAKTOR can process multiple disk volumes in one execution. Multiple disk volumes can be specified as a list of volume serials, a volume serial prefix, or all online disk volumes can be requested. The required online disks volumes are dynamically allocated.

CONTINUED . . .

40.03 CONTINUED . . .

**FAST
COMPAKTION
IN-PLACE**

If TYPE=FASTCPK is specified on the COMPAKT control statement, COMPAKTOR does a "COMPAKT-in-place" on one or more disk volumes, without taking or requiring a FDR backup. For each volume processed, COMPAKTOR analyzes the dumped VTOC and builds an internal map of the disk volume. A number of algorithms are then used to determine the "best" way of organizing the disk volume, and a new internal map is built; this new map is of the reorganized, or COMPAKTed, disk volume. User-specified data set placement or space release requests are honored during these algorithms.

COMPAKTOR moves tracks to their new locations according to the new volume organization map. Once this has been successfully done, COMPAKTOR rebuilds the extent information contained in the VTOC (and, if necessary, the VVDS).

Fast COMPAKTion achieves almost the same level of free space consolidation and reduction of multi-extent data sets as COMPAKTion-from-backup, without compromising the integrity of your data. FASTCPK maintains a log of its processing and provides an automatic recovery function which allows COMPAKTion to be restarted if the operating system or the job fails.

Fast COMPAKTion requires that there be at least a small amount of free space on the volume, in order to allocate a "recovery" data set used to keep track of COMPAKTOR's progress and allow restart if COMPAKTOR is prevented from completing its data movement.

Fast COMPAKTion supports most of the features of COMPAKTOR, except that you cannot change the size or location of the VTOC, VVDS or active indexed VTOC.

**SPACE
RELEASE**

If TYPE=RLSE is specified on the COMPAKT control statement, COMPAKTOR makes a special run whose only purpose is to release space from PS, PO, and ICF VSAM data sets. TYPE=RLSE does not move any data sets or attempt to reduce fragmentation or increase the size of free space. It does not require a backup of the disk volume; it simply reads the VTOC from the disk, manipulates it to release space as indicated by CPK options, and updates the VTOC in place. A TYPE=RLSE step will usually run in just a few seconds.

COMPAKTOR will print the name of each data set selected and the number of tracks released from each.

40.04 COMPAKTOR SPECIAL FEATURES**ICF VSAM
SUPPORT**

COMPAKTOR treats ICF VSAM clusters (including linear clusters and DB2 files) much like any other data set. Components will be relocated on the volume, and multi-extent components will usually be combined into a single extent (unless VSAMEXT=KEEP is specified, in which case CPK will attempt to position all the extents contiguously, one after the other). You may also request that CPK release space from over-allocated VSAM components, based on the HI-USED-RBA in the VVDS, by specifying VSRLSE=ALL on the CPK statement or on a SELECT statement that applies to the component. A percentage of the over-allocated space may be left as unused control areas for expansion by specifying %FREE= or %VSFREE= on the CPK statement or %FREE= on the SELECT statement.

CPK cannot release space from multi-volume VSAM components, and will only merge extents residing on the first volume of a multi-volume component. Also, CPK will not release space from any cluster where the "open for output" flag is on in the VVDS, even if the cluster is not currently allocated to another job, since the HI-USED-RBA may be inaccurate (IDCAMS VERIFY can reset that flag and correct the RBA).

By default, CPK will not relocate nor release space from components of ICF catalogs. To do so, you must specify OVERRIDE=YES on the CPK statement and provide a SELECT DSN= with the fully-qualified name of each catalog component (the catalog name for the data component, and the "catindex" name for the index component); space will not be released unless VSRLSE=ALL is specified on the SELECT statement.

CPK will never release space from a VVDS, and it cannot be moved by FASTCPK. CPK can merge extents and relocate the VVDS during a COMPAKT-from-backup if OVERRIDE=YES is specified and a SELECT selects it (the SELECT does not necessarily have to specify the fully-qualified VVDS name, a filter can select it).

It is the user's responsibility to ensure that the catalogs and VVDS are not in use when they are moved (the F CATALOG console command can be used to close them, see the appropriate IBM SYSTEM COMMANDS manual).

**COMPAKTING
ACTIVE
VOLUMES**

The data set enqueue option (DSNENQ=), coupled with the CPK unmovable data set table, gives you the capability to COMPAKT active volumes.

The DSNENQ= option requests that COMPAKTOR use a SYSDSN ENQ to determine which data sets on the volume being COMPAKTeD are currently active (allocated) and optionally ENQs inactive data sets to COMPAKTOR so that no other task can use them until COMPAKTOR is done with the volume. The COMPAKTOR unmovable table can be used to identify data sets which are not normally ENQed by the system but which should be considered unmovable anyways.

Space release (TYPE=RLSE) and Fast COMPAKTion (TYPE=FASTCPK) default to DSNENQ=USE and will automatically identify active data sets and bypass moving and releasing them.

COMPAKT-from-backup (TYPE=COMPAKT specified or defaulted) supports COMPAKTING active volumes only when the DUMP=YES option is specified (to invoke FDR for a backup. The DSNENQ= option must be specified on the FDR DUMP statement, not on the COMPAKT statement. Because of the elapsed time required, COMPAKT-from-backup is no longer recommended for active volumes; use FASTCPK.

Please read [Section 40.21](#) for details before attempting to COMPAKT any active volumes.

**FDR
INSTANTBACKUP**

If you are also licensed for FDR InstantBackup, it enhances FASTCPK to use the hardware features of certain disk subsystems to move data tracks from one volume to another without the channel overhead of sending the data to and from the CPU. This works on IBM RVA and StorageTek Iceberg and SVA systems with the Snapshot feature and EMC Symmetrix systems with the Timefinder feature. To invoke these hardware copy features, you must add the operand EMCCOPY=YES or SNAPSHOT=YES to the COMPAKT statement. The elapsed time of the FASTCPK will be significantly reduced. FDR InstantBackup for Symmetrix and RVA/Iceberg/SVA is described in [Sections 25](#) and [26](#).

CONTINUED . . .

40.04 CONTINUED . . .

**USED vs
ALLOCATED
TRACKS** Fast COMPAKTion by default only moves used tracks within PS, PO data sets (based on the "last block pointer") and ICF VSAM components (based on HI-USED-RBA) to reduce elapsed time, but you may specify DATA=ALL on the COMPAKT statement if you prefer to move all allocated tracks within all data sets.

CPK has special processing for PS and PO data sets which have a "last block pointer" of zeros; this may indicate an empty data set or one for which the LBP is not maintained. FASTCPK will move all allocated tracks of such data sets, and TYPE=RLSE will not release space from them.

**UNMOVABLE
DATA SETS** COMPAKTOR supports an internal table of unmovable data sets. You may specify data set names or groups that should always be considered unmovable. It is important that each installation update this table with the names of data sets which are active without being ENQed; see [Section 40.21](#) for a list of potential data sets for inclusion. This list of names is stored in a load module in the CPK program library, and can be built using either the FDR ISPF panels (option A.I.5) or program FDRZAPOP). Entries in the unmovable table and data sets identified as active can be moved by CPK SELECT statements if OVERRIDE=YES is specified on the CPK control statement. This option should be used with care. If OVERRIDE=NO is specified or defaulted, CPK will ignore any request to move or release space for any data set found in these tables.

**INDEXED
VTOC** COMPAKTOR can be used to reorganize a volume containing an active indexed VTOC. Fast COMPAKTion and space release will update the space maps in the VTOCIX directly, without disabling it. For COMPAKT-from-backup, COMPAKTOR will convert the volume with the indexed VTOC to an OS VTOC prior to COMPAKTion. After a successful COMPAKTion, CPK will then execute ICKDSF to rebuild the indexed VTOC. CPK will suppress the normal ICKDSF messages and operator replies.

CPK will not move the SYS1.VTOCIX data set unless you specify this data set name in the control statements and OVERRIDE=YES is indicated; you may wish to do this in order to move the VTOCIX next to the VTOC. Be sure that the volume is not online to any other CPU if the VTOCIX is moved.

**VOLUME
EXCLUSION** You may exclude (protect) certain volumes from COMPAKTOR processing by making special entries in the COMPAKTOR Unmovable Table. If COMPAKTOR finds a name in the Unmovable Table in the format:

```
FDRCPK . EXCLUDE . COMPAKT . VVVVVVV
```

any attempt to do any COMPAKTion on volume "vvvvvv" will result in an error. If the name is:

```
FDRCPK . EXCLUDE . RELEASE . VVVVVVV
```

then space release (TYPE=RLSE) will be inhibited. If the name in the table is a DSN entry, then only the named volume is excluded. If the entry is a DSG, then "vvvvvv" must be 5 or fewer characters and all volumes starting with that prefix will be excluded.

For example, entries of

```
DSN  FDRCPK . EXCLUDE . COMPAKT . VTS0123
DSG  FDRCPK . EXCLUDE . RELEASE . VTS0
```

will exclude volume TS0123 from COMPAKTion functions, and will exclude all TSO volumes from release processing.

You can use this volume exclusion to protect certain volumes or groups of volumes that should not normally be COMPAKTed against accidental COMPAKTion or Release.

CONTINUED . . .

40.04 CONTINUED . . .

**SMS
CONSIDERATIONS**

COMPAKTOR supports all functions on SMS-managed disk volumes.

Because of the restrictions of SMS, it is not possible to restore a SMS-managed volume to a new volume serial. If you are doing a COMPAKT-from-backup (COMPAKT TYPE=CPK) of a SMS-managed volume and you are restoring to a new volume, COMPAKTOR will force CPYVOLID=YES, so that the output disk will receive the volume serial of the original SMS volume. If the original volume is still online, the output volume will be forced offline.

NOTE: Innovation recommends using PGM=FDR to move entire volumes unless you have the need to move or expand the VTOC or move the VVDS.

COMPAKTOR supports all functions against these special types of data sets:

- PDSE (Partitioned Data Set Extended)
- HFS (Hierarchical File System) - used for storage of files for Unix System Services (Open Edition). They are treated like PDSEs by CPK.
- Extended Format including sequential striped data sets, sequential compressed data sets, and VSAM compressed KSDS clusters, supported only on SMS-managed volumes.
- Extended Addressability including VSAM clusters over 4GB in size, supported only on SMS-managed volumes.

**MODEL TYPE
CONVERSIONS**

COMPAKTOR may be used to convert a volume to a different capacity disk of the same device type. For example, COMPAKTOR may be used to convert a 3380 volume to another 3380 model, and to convert 3390 volumes to another 3390 model. COMPAKTOR automatically recognizes the number of cylinders on the output device and will update the VTOC and free space with the proper device size. This conversion can also be done with FDR, but COMPAKTOR offers several advantages:

During the conversion, COMPAKTOR can expand and reposition the VTOC, providing more room for data sets on a larger device. COMPAKTOR can also move the indexed VTOC and VVDS next to the new VTOC, plus any other data sets that should be close to the VTOC.

Unlike FDR, COMPAKTOR can convert a volume to a model with fewer cylinders than the original. In order for this to work, all allocated data sets on the original volume must fit in the space available on the output volume, and any unmovable data sets must be at addresses within the limits of the new device.

COMPAKTOR simulation supports model conversions, even if the new devices are not yet installed on your system, so that you can prepare for the actual conversion.

When several smaller disks are to be combined into one larger disk, COMPAKTOR may be used to convert one of the source disks, under its original volume serial, and FDRCOPY ([See Section 21](#)) may be used to add data sets from the other volumes.

COMPAKT-from-backup must be used for model conversions since Fast COMPAKTion (TYPE=FASTCPK) is always a "COMPAKT-in-place", back to the original disk.

40.05 COMPAKTOR JCL REQUIREMENTS

The running of COMPAKTOR requires the Job Control Language (JCL) statements detailed below.

**JOB
STATEMENT**

The JOB statement is user-specified and depends upon installation standards.

**EXEC
STATEMENT**

Must specify the COMPAKTOR program name – FDRCPK. May also contain region size required and PARM field data. The recommended region is a minimum of 2048K; a region of 4096K is required for a Fast COMPAKTion (TYPE=FASTCPK). You may use the PARM as the first control statement to specify the CPK, MAP, or SIM statement. Additional statements, if necessary, must be specified in the SYSIN data set.

**STEPLIB OR
JOBLIB DD
STATEMENT**

If required, must specify the load module library in which COMPAKTOR and FDR reside.

All libraries specified must be authorized.

**SYSPRINT DD
STATEMENT**

Specifies the output messages data set. This is a required DD statement and is usually a SYSOUT data set.

**SYSMAP DD
STATEMENT**

Specifies the output maps data set, usually a SYSOUT data set. If omitted, COMPAKTOR will dynamically allocate it to SYSOUT=*.

**SYSSUMM DD
STATEMENT**

Specifies the combined summary data set, and will contain a two-line summary of before/after statistics for each volume processed, sorted by volser; if omitted, the combined summary is printed at the end of the SYSMAP OUTPUT.

**SYSUDUMP DD
STATEMENT**

Specifies the ABEND dump data set. Although not required, we strongly urge you to always include this DD statement, so that we can help you diagnose error conditions. Usually specifies a SYSOUT data set.

**CPKWORK DD
STATEMENT**

Used for COMPAKT-from-backup if an active INDEXED VTOC exists on the volume being COMPAKted. Usually specifies a temporary disk data set (VIO is recommended). At least one track of space must be allocated. This data set must not be allocated on the volume being COMPAKted. It can be omitted for Fast COMPAKTion (TYPE=FASTCPK) and space release (TYPE=RLSE). If it is omitted, but COMPAKTOR requires it, it will be dynamically allocated as a 1-track temporary data set on UNIT=SYSALLDA.

**DISK1 DD
STATEMENT**

This DD statement must be omitted if you are using the VOL= operand to specify the disk volumes to be processed; since this is the recommended method for most CPK operations, the DISK1 DD is rarely used. If VOL= is used, CPK will fail if a DISK1 DD statement is found, since CPK will dynamically allocate the online disk volumes required to DISK1. The DISK1 DD is not used if you are doing a MAP or SIMULATE from an FDR backup tape.

If specified, it must include UNIT=, VOL=SER=, and DISP=OLD or SHR for the single disk volume to be processed.

EXAMPLE: //DISK1 DD UNIT=3390,VOL=SER=ABC123,DISP=OLD
The DD name must be DISK1 if DUMP=YES is specified for a COMPAKT-from-backup since FDR will be invoked to backup DISK1 to TAPE1; for other functions you can specify any DD name, but you must indicate the name on a FROMDD= or TODD= operand on a CPK control statement.

CONTINUED . . .

40.05 CONTINUED . . .

TAPE1 DD STATEMENT For a COMPAKT-from-backup (TYPE=COMPAKT specified or defaulted), specifies the input data set containing the FDR backup which will be used as input to COMPAKTOR. If the backup is cataloged, only DSN= and DISP=OLD are required.

If DUMP=YES is specified, TAPE1 specifies the output data set to which FDR will backup the disk volume and usually requests scratch tapes. You probably want to catalog this backup so that CPK can be easily restarted in case of a failure. Following the backup, CPK will read this backup for the COMPAKTion. Do not specify FREE=CLOSE.

If mapping or simulating from tape, TAPE1 specifies the first tape volume of an FDR dump. Volumes other than the first need not be specified.

The DD name need not be TAPE1 unless DUMP=YES is specified. However, if it is not, you must specify the name in the FROMDD= operand on a major control statement.

For Fast COMPAKTion (TYPE=FASTCPK) and RELEASE (TYPE=RLSE), this DD statement is not used and should be omitted.

TAPE11 DD STATEMENT If COMPAKTing with DUMP=YES, specifies the optional duplicate tape to be created by FDR during the dump. Under MVS, the JCL option FREE=CLOSE may be used to deallocate this tape drive before CPK begins reading the TAPE1 backup. Do not specify DISP=(NEW,CATLG) and FREE=CLOSE together for TAPE11 if this data set is to be cataloged in an ICF VSAM catalog residing on the volume being COMPAKTed.

This DD statement is never required for COMPAKTOR and is ignored on any operation except COMPAKT-from-backup with DUMP=YES.

SYSIN DD STATEMENT Specifies the control statements data set. Usually a SYSIN data set. If omitted, COMPAKTOR assumes that you wish to simulate a reorganization of the disk volume specified by the DISK1 DD statement (SIMULATE TYPE=COMPAKT).

40.06 COMPAKTOR STATEMENTS

COMPAKTOR statements are of two types: major statements and minor statements.

**MAJOR
STATEMENTS**

Only a single major statement may be used and it must be the first statement input. The major statement determines the COMPAKTOR run type: mapping, simulation or COMPAKTion. The following major statements are supported:

- MAP** – This statement tells COMPAKTOR to map one or more disk volumes, or one FDR backup of a disk volume.
- SIMULATE** – This statement causes COMPAKTOR to simulate reorganization of one or more disk volumes, or one FDR backup of a disk volume.
- COMPAKT** – This statement causes COMPAKTOR to reorganize one or more disk volumes.
- If TYPE=RLSE is specified on a SIMULATE or COMPAKT statement, CPK releases unused space within PS, PO, and ICF VSAM data sets.

If no major statement is input to COMPAKTOR, it assumes a default of SIMULATE.

**MINOR
STATEMENTS**

Minor statements must be placed after the major statement, and may be coded as often as required. The following minor statements are supported:

- SELECT** – This statement may be used to identify a data set or group of data sets which require special handling by COMPAKTOR.
- SEQUENCE** – This statement may be used to define a set of data sets to be placed in a prescribed sequence (a "sequenced set").
- ENDSEQ** – This statement must be used if a SEQUENCE statement was specified, to denote the end of the sequenced set. One or more SELECT statements must appear between the SEQUENCE and ENDSEQ statements.

When SELECT DSN=filter statements are used, they must follow any SEQUENCE, ENDSEQ and SELECT DSN= statements specifying fully-qualified names.

**FDR
STATEMENTS**

In addition to the COMPAKTOR statements, CPK will accept the FDR DUMP control statement. When you are doing a COMPAKT-from-backup (TYPE=COMPAKT) and you specify DUMP=YES to cause CPK to invoke FDR to dump the disk volume, you may want to specify options for FDR to use during the dump. If so, include a DUMP TYPE=FDR statement preceding the COMPAKTOR statements. This statement, if used, must be the first statement in the SYSIN data set, and must be contained on one record (i.e. must not be continued). Some of the other operands that are valid are COMPRESS=, MAXERR= and DSNENQ=. Do not specify the ENQ= or ENQERR= operands; they will be ignored. These operands are described in [Section 10.04](#).

EXAMPLE : DUMP TYPE=FDR , DSNENQ=USE
 COMPAKT DUMP=YES

NOTE: Innovation recommends using FASTCPK (CPK TYPE=FASTCPK) for most COMPAKTions; it does not invoke FDR and does not require the DUMP statement.

DEFAULTS

The defaults for certain operands, such as ACTMESS=, LINECNT=, SIZEKEEP=, NOSECOND=, and HILIGHT= may be permanently changed in the FDR option table. These may be updated via the FDR ISPF support or by program FDRZAPOP. Defaults shown in the descriptions which follow are the defaults distributed with COMPAKTOR.

CONTINUED . . .

40.07 MAP STATEMENT

MAP

ENQ=YES|NO|RESERVE

,FROMDD=ddname

,HIGHLIGHT=YES|NO

,LINECNT=nn

,MAPFORMAT=NEW|OLD

,MAPS=EXTENTS|FREESPC|SUMMARY

,STORGRP=(storgrp,...)

,VOL=(volser,volser,...)

MAP STATEMENT Th MAP statement is used to generate a map of disk volumes. You may map either one or more online disk volumes or the backup of one disk volume from a FDR full-volume backup data set. Full VTOC error checking will be performed in either case. The input device, whether disk or FDR backup, is never altered in any way.

OPERANDS

ENQ=

YES – specifies that CPK will enqueue the VTOC for the duration of the MAP (but not RESERVE it).

NO – specifies that CPK will not enqueue the VTOC.

RESERVE – specifies that CPK will issue a ENQ/RESERVE against the VTOC on the volume being mapped.

CAUTION: If ENQ=YES is coded on a shared DASD system or if ENQ=NO is coded on any system, CPK may produce incorrect results or go into a loop if data sets are allocated or scratched during the MAP. Since the RESERVE is short duration for a MAP, Innovation does not recommend overriding the default.

FROMDD=

Default is RESERVE. Ignored if mapping an FDR backup.

Specifies the dd-name of the DD statement pointing to the input device. In order to map a FDR backup residing on a disk backup file (not tape), you must specify FROMDD=TAPEX, otherwise, CPK will map the volume on which the backup resides instead.

HIGHLIGHT=

HI

Default is DISK1.

YES – specifies that CPK will highlight certain data on the maps in order to make it stand out. This is done by overprinting the line three times.

NO – specifies no highlighting.

NOTE: When formatting for a display terminal or a laser printer, we suggest you specify NO highlighting.

LINECNT=

LC

Default is YES unless overridden in the FDR option table.

Specifies the number of lines COMPAKTOR is to print per page, from 10 to 999.

Default is 58 unless overridden in the FDR option table.

CONTINUED . . .

40.07 CONTINUED . . .

MAPFORMAT= NEW – specifies that a new format for the "before" and "after" maps will be used, which supports 5-digit cylinder numbers to accommodate very large disks such as the 3390-9.

OLD – specifies that the old map format, with 4-digit cylinder numbers and 5-digit relative track numbers, is to be used. If MAPFORMAT=OLD is specified for the large disks, inaccurate cylinder and track numbers may be printed.

Default is NEW for disks with over 65535 tracks (such as the 3390-9) and OLD for all other disks.

MAPS= Specifies which maps are to be printed on SYSMAP.

EXTENTS – specifies that a full track/extent map plus a summary report is desired.

FREESPC – specifies that CPK is to print only free areas.

SUMMARY – specifies that only a summary report will be produced.

Default is EXTENTS.

STORGRP= Specifies the names of one or more SMS storage groups. COMPAKTOR will interface to SMS to obtain all of the disk volume serials which are defined to each of the specified storage group names, and will process them as if a VOL= operand were specified for them.

The syntax for STORGRP= is the same as for VOL= as documented below. You may specify a single storage group, a list of groups, and/or a "group" of group names (e.g., STORGRP=DB*). You cannot specify STORGRP= and VOL= in the same run.

CONTINUED . . .

40.07 CONTINUED . . .

VOL=

Specifies the disk volume serial(s) which COMPAKTOR is to dynamically allocate and map. FROMDD= and VOL= may not both be specified. If VOL= is specified, there must not be a DISK1 DD statement in the COMPAKTOR JCL.

A single volume serial may be specified as VOL=volser or multiple volume serials may be specified:

- 1) A list of volume serials may be given, enclosed in parentheses, up to a maximum of 254, e.g.,

VOL=(TSO001,TSO002,TSO003)

COMPAKTOR will map the volumes in the order specified.

- 2) A volume group may be specified by placing an asterisk at the end of the volser prefix, e.g.,

VOL=TSO*

Within that group, COMPAKTOR will map the volumes in the order that their UCBs exist in the MVS system.

- 3) The two may be combined, e.g.,

VOL=(TSO*,PROD*,ABC001)

Each group will be mapped in the order described above before going on to the next group or volume.

- 4) All online disk volumes may be selected by:

VOL=*

COMPAKTOR will map the volumes in the order that their UCBs exist in the MVS system.

COMPAKTOR will not check for duplicate volume serials, so volumes selected more than once will be mapped more than once.

40.08 COMPAKT AND SIMULATE STATEMENTS

This page lists the operands available for Fast COMPAKTion (TYPE=FASTCPK), a COMPAKT-in-place that does not use a backup. Only the major statement (COMPAKT or SIMULATE) and TYPE=FASTCPK are required, all other operands are optional.

COMPAKT CPK SIMULATE SIM	TYPE=FASTCPK FCPK	,NOSECOND=RLSE NORLSE
	,ACTMESS=YES <u>NO</u>	,OVERRIDE=YES <u>NO</u>
	,CONFMESS=YES <u>NO</u>	,PORLSE=ALL ROUND TRK ,PSRLSE=ALL ROUND TRK
	,CPKFRAGI=nnn ,CPKFREEX=nn	,PRINT=ALL
	,DATA=ALL <u>USED</u>	,SELTERR= <u>YES</u> <u>NO</u>
	,DSNENQ=NONE TEST USE	,SIZEKEEP=(size,pct,mx)
	,EMCCOPY=YES <u>NO</u>	,SNAPSHOT=YES <u>NO</u>
	,ENQ=YES <u>NO</u> <u>RESERVE</u>	,STORGRP=(storgpr,...) ,VOL=(volser,volser,...)
	,EXTENTS=KEEP	,UNABLE=IGNORE
	,HILIGHT=YES <u>NO</u>	,VSRLSE=ALL <u>NO</u>
	,LBPZERO=VALID INVALID	,VSAMEXT=KEEP <u>MERGE</u>
	,LINECNT=nn	
	,LOG=YES <u>NO</u>	,%FREE=nn ,%PSFREE=nn ,%POFREE=nn ,%VSFREE=nn
	,MAPFORMAT=OLD NEW	
	,MAPS=(<u>ALL</u> AFT SUMMARY, <u>EXTENTS</u> FREESPC)	

Certain options that were documented in previous releases, such as BADCHAIN=, CPKDSNMX=, CPKMULTX=, FIT=, LBPZERO=, LRDAY=, MINRLSE=, OBJECT=, and SORT= have been removed from this manual because they are rarely used or not recommended. For documentation on these options, see the member \$\$CPK in the FDR ICL (Installation Control Library).

CONTINUED . . .

40.08 CONTINUED . . .

This page lists the operands available for COMPAKTOR RELEASE (TYPE=RLSE), the fast space release option. Only the major statement (COMPAKT or SIMULATE) and TYPE=RELEASE are required, all other operands are optional.

COMPAKT CPK SIMULATE SIM	TYPE=RELEASE RLSE	,NOSECOND=RLSE NORLSE
	,ACTMESS=YES <u>NO</u>	,OVERRIDE=YES <u>NO</u>
	,CONFMESS=YES <u>NO</u>	PORLSE=ALL ROUND TRK ,PSRLSE=ALL ROUND TRK
	,DSNENQ=NONE TEST USE	,PRINT=ALL
	,ENQ=YES NO <u>RESERVE</u>	,SELTERR= <u>YES</u> NO
	,HILIGHT=YES NO	
	,LBPZERO=VALID INVALID	,STORGRP=(storgrp,...) ,VOL=(volser,volser,...)
	,LINECNT=nn	,VSRLSE=ALL <u>NO</u>
	,LOG=YES <u>NO</u>	,%FREE=nn ,%PSFREE=nn ,%POFREE=nn ,%VSFREE=nn
	,MAPFORMAT=OLD NEW	
	,MAPS=(ALL AFT <u>SUMMARY</u> ,EXTENTS FREESPC)	

Certain options that were documented in previous releases, such as BADCHAIN= and MINRLSE= have been removed from this manual because they are rarely used or not recommended. For documentation on these options, see the member \$\$CPK in the FDR ICL (Installation Control Library).

CONTINUED . . .

40.08 CONTINUED . . .

This page lists operands available for a COMPAKT-from-backup (TYPE=COMPAKT). Only the major statement name (COMPAKT or SIMULATE) is required; all operands are optional. If TYPE= omitted a COMPAKT-from-backup will be done (see the preceding pages for required operands for FASTCPK and RELEASE).

COMPAKT CPK SIMULATE SIM	TYPE= <u>COMPAKT</u> CPK	
,ACTMESS=YES NO	,LOG=YES <u>NO</u>	,SMSPROT= <u>ALL</u> NONE
,CONFMESS= <u>YES</u> NO	,MAPFORMAT=NEW OLD	,TEMPS=KEEP
,CPKFRAGI=nnnn ,CPKFREEEX=nn	,MAPS=(<u>ALL</u> AFT SUMMARY, EXTENTS FREESPC)	,TODD=ddname
,CPYVOLID= <u>YES</u> NO	,NOSECOND=RLSE NORLSE	,VOL=volser
,DUMP=YES <u>NO</u>	,OVERRIDE=YES <u>NO</u>	,VSRLSE= <u>ALL</u> <u>NO</u>
,ENQ=YES NO <u>RESERVE</u>	,PORLSE=ALL ROUND TRK ,PSRLSE=ALL ROUND TRK	,VSAMEXT=KEEP <u>MERGE</u>
,EXTENTS=KEEP	,PRINT=ALL	,VTOC= <u>NOCHANGE</u> COMPAKT
,FROMDD=ddname	,SELTERR= <u>YES</u> NO	,%FREE=nn ,%PSFREE=nn
,HIGHLIGHT=YES NO	,SIZEKEEP=(size,pct,mx)	,%POFREE=nn ,%VSFREE=nn
,LINECNT=nn		

Certain options that were documented in previous releases, such as BADCHAIN=, CPKDSNMX=, CPKMULTX=, DEVCYL=, FIT=, LBPZERO=, LRDAY=, MINRLSE=, OBJECT=, and SORT= have been removed from this manual because they are rarely used or not recommended. For documentation on these options, see the member \$\$CPK in the FDR ICL (Installation Control Library).

Innovation strongly recommends the use of Fast COMPAKTion (COMPAKT TYPE=FASTCPK) instead of COMPAKTion-from-backup (COMPAKT TYPE=CPK). FASTCPK does not use an FDR backup and will run much faster than COMPAKT-from-backup, typically just minutes to reorganize a volume.

CONTINUED . . .

40.08 CONTINUED . . .

COMPAKT STATEMENT	The COMPAKT or CPK statement invokes either a full COMPAKTion of one or more disk volumes, or causes just release of free space in data sets, depending on the TYPE= operand.
SIMULATE STATEMENT	The SIMULATE or SIM statement invokes a simulation of a real COMPAKTion, so that you can see what the results would be if the COMPAKTion was actually done. The before and after maps and summary (depending on the MAPS= operand) will be printed. TYPE=RLSE can be specified to simulate space release. Note that a SIMULATE TYPE=COMPAKT will not test for active data sets so when COMPAKTing an active volume the simulation results may differ from the actual results unless you specify DSNENQ=USE on the SIMULATE statement. However, TYPE=RLSE and TYPE=FASTCPK will test for active data sets and act accordingly. Simulation may be done against the actual disk volume, or, for TYPE=COMPAKT, against an FDR full-volume backup data set.
OPERANDS	Abbreviations for the operands are shown under the operand name where applicable.

TYPE=	<p>FASTCPK or FCPK – specifies that COMPAKTOR is to perform a Fast COMPAKTion-in-place on one or more disk volumes. No FDR backup is needed or used. TYPE=FASTCPK assumes that the volume may be active and assumes ACTMESS=NO, CONFMESS=NO, TEMPS=KEEP and DSNENQ=USE (NONE for SIMULATE); active data sets will not be moved. Data sets may not be scratched, the VTOC cannot be moved or expanded, and the VVDS cannot be moved.</p> <p>RELEASE or RLSE – specifies that COMPAKTOR is to release excess space from PS, PO and/or VSAM data sets on one or more disk volumes. No FDR backup is needed or used. At least one operand requesting space release must be specified on the COMPAKT/SIMULATE statement or on a SELECT statement. TYPE=RLSE assumes that the volume may be active and assumes ACTMESS=NO, CONFMESS=NO, DSNENQ=USE (NONE for SIMULATE); space will not be released from active data sets. No data sets will be moved.</p> <p>COMPAKT or CPK – specifies that COMPAKTOR is to perform a COMPAKTion on one disk volume, from a FDR full-volume backup. If DUMP=YES is specified, COMPAKTOR will invoke FDR to take that backup and then read it back. Otherwise, the backup must have been taken previously. TYPE=COMPAKT assumes that the volume is not active and that all data sets are movable; FASTCPK should be used to COMPAKT active volumes. You must use TYPE=COMPAKT instead of FASTCPK if the output volume is not the same disk device as the input volume, if data sets are to be scratched, or if the VTOC is to be moved or expanded or if the VVDS is to be moved.</p> <p>Default is TYPE=COMPAKT. Innovation strongly recommends TYPE=FASTCPK instead of the default of TYPE=COMPAKT.</p>
ACTMESS=	<p>YES – specifies that CPK will notify the operator if the volume being COMPAKTed is currently active (has more than one open DCB). The operator must then reply to the FDRW81 message before CPK will continue COMPAKTing the volume.</p> <p>NO – specifies that CPK will suppress the active message.</p> <p>Default is YES for TYPE=COMPAKT unless overridden in the FDR option table. Default is NO for TYPE=RLSE and TYPE=FASTCPK. ACTMESS is ignored for SIMULATE or VTOC=COMPAKT.</p>

CONTINUED . . .

40.08 CONTINUED . . .

CONFMESS=	<p>YES – specifies that CPK will ask the operator, via a WTOR, if the COMPAKTion can proceed. The operator must reply YES to the FDRW80 message to begin the COMPAKTion.</p> <p>NO – specifies that CPK begin the COMPAKTion without asking the operator for confirmation.</p> <p>Default is YES for TYPE=COMPAKT and NO for TYPE=RLSE and TYPE=FASTCPK. CONFMESS is ignored for SIMULATE.</p>
CPKFRAGI= CPKFREEEX=	<p>These operands define an amount of fragmentation that you consider acceptable. If neither of these operands has a non-zero value, COMPAKTion will be done unconditionally. If one or both of these operands is non-zero, COMPAKTion of a given volume will be performed only if one of them is exceeded on that volume. These operands can only be specified with TYPE=COMPAKT,DUMP=YES or TYPE=FASTCPK.</p> <p>CPKFRAGI – specifies the largest fragmentation index that a volume may have before it is to be COMPAKTed. The index is calculated from an IBM formula and is discussed in Section 40.15. A decimal point is assumed in front of the value, e.g., CPKFRAGI=2, CPKFRAGI=20 and CPKFRAGI=200 all refer to an index of .200. Innovation recommends the use of CPKFREEEX= instead of CPKFRAGI= since the IBM fragmentation index does not always reliably indicate the state of the volume.</p> <p>CPKFREEEX – specifies the maximum number of free space areas that can exist before the volume is to be COMPAKTed.</p>
CPYVOLID= CPY	<p>NO – specifies that the volume serial number of the target volume (the volume specified in the TODD= parameter) is to be retained. This applies only to TYPE=COMPAKT when the target volume is not the same as the original volume.</p> <p>YES – the target volume receives the serial number of the dumped volume.</p> <p>The default is YES.</p>
DATA=	<p>ALL – specifies that all allocated tracks of PS, PO and ICF VSAM data sets will be moved if the data set is moved by a Fast COMPAKTion, except for tracks that are released by COMPAKTOR.</p> <p>USED – specifies that only the used tracks of PS, PO and ICF VSAM data sets will be moved by a Fast COMPAKTion (TYPE=FASTCPK). This will improve COMPAKTOR elapsed time and should be used unless there are data sets on the volume whose last block pointer or HI-USED-RBA is unreliable.</p> <p>Default is USED.</p>

CONTINUED . . .

40.08 CONTINUED . . .

DSNENQ=

Specifies that COMPAKTOR is to enqueue all of the data sets on the volume being processed, issuing an exclusive enqueue with a major name of 'SYSDSN' and a minor name of the data set name. This is the enqueue used by the job scheduler. If the enqueue fails, COMPAKTOR will issue a warning message for the data set and will consider it to be an active data set (the data set will not be moved and will not have any space released). Data sets which are not moved or released will not be enqueued (for Fast COMPAKTion they may be enqueued briefly and dequeued).

DSNENQ= is independent of ENQ= which prevents new data sets from being allocated and old data sets from being scratched.

DSNENQ= can be used only if TYPE=RLSE or TYPE=FASTCPK has been specified.

TEST – The data sets will only be tested to see if they are active at the time the COMPAKTion starts. The data sets will not be enqueued.

USE – The data sets will be enqueued for the duration of the COMPAKTion of the volume. If they are enqueued to another task, they will be considered active.

NONE – No data set ENQ will be issued; active data sets will not be detected and may be moved or released.

NOTE: If the data set is specified in a DD statement in the CPK job with DISP=SHR, the scheduler enqueue for the data set will be changed to EXCLUSIVE (DISP=OLD) if you specify TEST or USE.

Default is USE for COMPAKT and NONE for SIMULATE, but only for the TYPE=FASTCPK and TYPE=RLSE options. For COMPAKT TYPE=COMPAKT with DUMP=YES, the DSNENQ= option on the FDR DUMP statement **must be used** to detect active data sets.

CAUTION: This option may not be effective on shared DASD unless a cross-system enqueue facility (such as GRS or CA-MIM) is available and the SYSDSN QNAME is broadcast across systems. Without this capability, COMPAKTOR can only determine which data sets are active on the system it is running on.

For further information on ENQs, see the member ENQ in the FDR ICL Library.

**DUMP=
D**

YES – specifies that COMPAKTOR must invoke FDR to dump the disk volume prior to COMPAKTion. The dump is always from the DISK1 DD statement to the TAPE1 statement; this is not affected by the FROMDD= or TODD= operands. Note that both DD statements are required if DUMP=YES is coded. However, DISK1 may be omitted if VOL= is specified.

NO – CPK is to restore from an existing FDR backup tape.

NOTE: DUMP= applies only to TYPE=COMPAKT and is ignored for SIMULATE.

The default is NO.

CONTINUED . . .

40.08 CONTINUED . . .

- EMCCOPY=** **YES** - Enables the use of a hardware feature (EMCCOPY) on EMC Symmetrix disk systems to move tracks internally without sending the data to and from the CPU, resulting in a reduction in elapsed time. EMCCOPY=YES can be specified only if you are also licensed for FDR InstantBackup (see [Section 25](#)). When a FASTCPK detects that the disk volume are in a Symmetrix which supports EMCCOPY, it is automatically invoked.
- NO** - EMCCOPY will not be used, even if available. Data tracks will be moved by reading and writing them.
- ENQ=** **YES** – specifies that CPK will enqueue the VTOC for the duration of the operation (but not RESERVE it).
- NO** – specifies that CPK will not enqueue the VTOC. ENQ=NO is not valid for TYPE=FASTCPK and TYPE=RELEASE.
- RESERVE** – specifies that CPK will issue a ENQ/RESERVE against the VTOC on the volume being COMPAKTed. The volume will not be accessible by other systems sharing this DASD volume until the COMPAKTion completes, unless your cross-CPU ENQ system converts SYSVTOC RESERVEs into cross-CPU ENQs.
- NOTE:** Additional information is available on cross system ENQ considerations, in the FDR Installation Control Library (ICL) under the member name ENQ.
- CAUTION: If ENQ=YES is coded on a shared DASD system or if ENQ=NO is coded on any system, CPK may lose data sets or go into a loop if data sets are allocated or scratched during COMPAKTion.*
- Default is RESERVE.
- For further information on ENQs, see the member ENQ in the FDR ICL Library.
- EXTENTS=** **KEEP** – specifies that all multi-extent non-VSAM data sets are to retain the number and size of their extents. If used with PSRLSE= and/or PORLSE=, the unused space will be freed down to the requested boundary, but the extents occupied by the remaining used portion of the data set will not be merged.
- If omitted, COMPAKTOR automatically attempts to reduce most multi-extent data sets down to a single extent. This does not apply to TYPE=RLSE, since remaining extents are not moved.
- FROMDD=** Specifies the ddname of the input DD statement.
- FDD**
- For TYPE=FASTCPK or TYPE=RLSE it points to the disk volume to be processed.
 - For COMPAKT TYPE=COMPAKT this must point to an FDR backup file on tape or disk.
 - For SIMULATE TYPE=COMPAKT, if you wish to simulation COMPAKTion from a FDR backup data set on disk, you must specify FROMDD=TAPEX. If the FROMDD= does not start with TAPE, COMPAKTOR will simulate COMPAKTion of the disk volume pointed to by the DD statement.
- Default is TAPE1 on COMPAKT TYPE=COMPAKT, DISK1 in all other cases.

CONTINUED . . .

40.08 CONTINUED . . .

HILIGHT=	YES – specifies that CPK will highlight certain data on the maps in order to make it stand out. This is done by overprinting the line three times.
HI	NO – specifies no highlighting.
	NOTE: when formatting for a display terminal or a laser printer, we suggest you specify NO highlighting.
	Default is YES unless HILIGHT=NO is set in the FDR option table.
LINECNT=	Specifies the number of lines COMPAKTOR is to print per page, from 10 to 999.
LC	Default is 58 lines to a page unless overridden in the FDR option table.
LOG=	YES – specifies that CPK will issue console operator messages (FDRW82) indicating when begins and ends modification of a disk volume. This may aid the operator in determining whether a COMPAKTOR operation has completed successfully. Innovation recommends LOG=YES so that operators are aware that COMPAKTions are in progress.
	NO – specifies that no FDRW82 console messages will be issued.
	Default is NO.
MAPFORMAT=	NEW – specifies that a new format for the "before" and "after" maps will be used, which supports 5-digit cylinder numbers to accommodate very large disks such as the 3390-9.
	OLD – specifies that the old map format, with 4-digit cylinder numbers and 5-digit relative track numbers, is to be used. If MAPFORMAT=OLD is specified for the large disks, inaccurate cylinder and track numbers may be printed.
	Default is NEW for disks with over 65535 tracks (such as the 3390-9) and OLD for all other disks.
MAPS=	Specifies which maps are to be printed on SYSMAP:
	ALL- specifies that both the before and after maps are to be printed.
	AFT- specifies that only the 'after' COMPAKTion/simulation map is to be printed.
	EXTENTS- specifies that a full track/extent map is desired.
	FREESPC- specifies that CPK is to print only free areas.
	SUMMARY- specifies that only a summary report will be produced (no detail maps).
	One or more of these may be entered in any order; if more than one is used, they must be surrounded by parenthesis.
	Default is (ALL,EXTENTS) except for TYPE=RLSE where the default is SUMMARY.

CONTINUED . . .

40.08 CONTINUED . . .

NOSECOND=	<p>Specifies how COMPAKTOR will treat PS, PO and ICF VSAM data sets which do not have any secondary allocation quantity defined.</p> <p>NORLSE – do not release space from any data set which does not have a secondary allocation.</p> <p>RLSE – data sets with no secondary allocation are candidates for space release.</p> <p>The default is RLSE unless overridden in the FDR option table (See Section 91 or 92).</p>
OVERRIDE=	<p>YES – specifies that CPK will honor SELECT control statements to move or release data sets specified in the unmovable table, or which were flagged as "active" by the DSNENQ= option, or other types as specified in section 40.12.</p> <p>NO – specifies that CPK will not move or release data sets specified in the unmovable table or which are active or are unmovable for other reasons.</p> <p><i>WARNING: This option should be used with care, since it may allow data sets that are in use to be moved, causing system or application failures. See Section 40.21 for details on COMPAKting active volumes.</i></p> <p>Default is NO.</p>
PSRLSE= PSR	<p>Specifies that unused tracks within physical sequential (DSORG=PS) data sets are to be freed (this includes extended format (striped) PS data sets). CPK uses the last block pointer (DS1LSTAR) field in the VTOC to determine how many tracks of the data set are used. If this field is all zeroes, CPK will not free any space from the data set. Some additional tracks may be retained as free space for expansion if %FREE= or %PSFREE= is coded.</p> <p>ALL – for track-allocated data sets, all unused tracks are to be freed. For cylinder-allocated data sets, only unused cylinders are freed.</p> <p>ROUND – only wholly unused extents are to be freed.</p> <p>TRK – all of the unused tracks are to be freed even for data sets which are cylinder allocated.</p> <p>NOTE: If PSRLSE=TRK releases a cylinder-allocated data set past its cylinder boundary, CPK will reset the cylinder-allocated indication in the affected extent descriptors in the DSCB. However, the secondary allocation type will remain the same as before the release. (The secondary allocation type is shown in the SPACE ALLOC field in the CPK map). Some products, such as ISPF, test the secondary allocation type and may display a data set as occupying a number of whole cylinders when it really does not.</p> <p>If omitted, all PS data sets retain their original number of tracks unless they are affected by the RLSE= operand of a SELECT statement.</p>
PORLSE= POR	<p>This parameter is identical to the PSRLSE parameter, but applies only to partitioned data sets (DSORG=PO). This also applies to PDSE (partitioned extended) data sets, except that an internal high-block number is used for the free space calculation; space will not be released from PDSEs by a COMPAKT-from-backup without DUMP=YES. %FREE= or %POFREE= may cause some free space to be retained for expansion.</p>
PRINT=ALL	<p>Specifies that CPK is to print a list of the unmovable and active data sets.</p>

CONTINUED . . .

40.08 CONTINUED . . .

SELTERR= **YES** – specifies that CPK will terminate processing on a volume if:

- a SELECT statement is entered which does not match any data set on the volume.
- a SELECT statement is entered for an unmovable data set (either from the COMPAKTOR unmovable table or identified as active by the DSNENQ= option), unless that SELECT specifies POS=KEEP.

NO – specifies that CPK will ignore SELECT statements meeting the above conditions.

Default is YES unless SELTERR=NO is set in the FDR option table.

SIZEKEEP= This is a performance option which reduces the number and size of the data sets to be moved by COMPAKTOR while still providing most of the free space consolidation benefits of COMPAKTion. Since data sets which are not moved will not have their tracks restored, this will significantly reduce COMPAKTOR elapsed time. There are 3 subparameters:

The first subparameter "size" specifies that COMPAKTOR is to search for groups of allocated tracks greater than or equal to this size; a group is a contiguous set of tracks, belonging to one or more data sets, with a free space extent on either side of it. These groups will be initially made unmovable. "size" may range from 0 to 999999.

The second parameter "pct" specifies a percentage reduction in the number of free space areas that COMPAKTOR must achieve (e.g., 50 specifies that CPK must reduce the number of free space areas on the volume by at least 50 percent). If COMPAKTOR cannot achieve goal, it will make some of the track groups selected by the first subparameter movable (the smaller groups in the list) and try again, repeating this until the desired percentage reduction is achieved or until all groups have been made movable. "pct" may range from 1 to 100.

The third parameter "mx" is the maximum number of extents that a multi-extent data set may have to be included in SIZEKEEP processing. Data sets with more than "mx" extents will not be considered part of any SIZEKEEP track group, and COMPAKTOR will try to combine their extents. If some, but not all, of the extents of a data set with "mx" or fewer extents are part of SIZEKEEP track groups, COMPAKTOR may move the other extents but will not combine them. "mx" may range from 1 to 60.

If a data set which would be made unmovable by SIZEKEEP is positioned by a SELECT statement with POS=, the SELECT will be honored. Space release is honored for SIZEKEEP data sets. SIZEKEEP is honored only for COMPAKTion in-place (TYPE=FASTCPK or TYPE=COMPAKT).

Specifying SIZEKEEP=0 disables this option for a COMPAKT-from-backup, but Fast COMPAKTion may still apply it to groups over 100 tracks as long as it can reduce the volume to 1 or 2 free areas.

The default is SIZEKEEP=(100, 90, 60). These defaults can be overridden in the FDR Global Option Table. The parenthesis can be omitted if only "size" is specified. For detailed recommendations on the use of SIZEKEEP= see the COMPAKTOR HINTS AND RECOMMENDATIONS in [Section 40.25](#).

Innovation strongly recommends that you run with the default SIZEKEEP= option.

CONTINUED . . .

40.08 CONTINUED . . .

- SMSPROT=** **ALL** – specifies that CPK will force CPYVOLID=YES when the input volume has the SMS-managed indicator in the VTOC, and will enforce the rule that a SMS-managed volume can only be restored to a volume whose UCB indicates SMS-management, and that a non-SMS volume must be restored to a non-SMS volume. This applies only to TYPE=COMPAKT when the target volume is not the original volume.
- NONE** – bypasses the above SMS checks, and honors CPYVOLID=NO if specified.
- Default is ALL.
- TEMPS=** **KEEP** – specifies that temporary data sets are to be retained when
T COMPAKTING. This is forced for TYPE=RLSE and TYPE=FASTCPK.
- Default is that temporary data sets are automatically deleted under TYPE=COMPAKT.
- TODD=** For COMPAKT TYPE=COMPAKT, specifies the ddname of the DD statement pointing to the target disk volume of the COMPAKTion. Must be a direct-access device.
- Default is DISK1.
- NOTE:** If DUMP=YES is coded, it is possible to direct CPK to write the compacted disk volume to a different disk than the one that was dumped. The disk pointed to by DISK1 will be dumped to TAPE1 by FDR, but CPK will modify the volume specified by the DD statement pointed to by TODD=; that DD name must NOT start with "DISK". This may be used to move a disk volume to a new device and COMPAKT it or reposition the VTOC at the same time.
- For SIMULATE TYPE=COMPAKT, may specify the dd-name of a DD statement that points to a direct-access device with the capacity (i.e. number of cylinders) desired for the target volume. Must be a device of the same basic type as the source volume. For example, if the source volume is a 3380 of any capacity (single,double, or triple), then TODD may specify a 3380 of any capacity, but not a 3390. The TODD volume will not be altered in any way by SIMULATE.
- The default is that the target device is considered to have the same capacity (same number of cylinders) as the source volume.
- For TYPE=RLSE or TYPE=FASTCPK, this keyword must not be used.
- SNAPSHOT=** **YES** - Enables the use of Snapshot on IBM RVA or StorageTek Iceberg or SVA disk systems with the Snapshot feature to "instantly" move data tracks from one location to another, resulting in a significant reduction in elapsed time. SNAPSHOT=YES can be specified only if you are also licensed for FDR InstantBackup ([see Section 26](#)). When a FASTCPK detects that the disk volume is in a RVA/Iceberg/SVA which supports Snapshot, it is automatically invoked.
- NO** - Snapshot will not be used, even if available. Data tracks will be moved by reading and writing them.

CONTINUED . . .

40.08 CONTINUED . . .

STORGRP= Specifies the names of one or more SMS storage groups. COMPAKTOR will interface to SMS to obtain all of the disk volume serials which are defined to each of the specified storage group names, and will process them as if a VOL= operand were specified for them.

The syntax for STORGRP= is the same as for VOL= as documented below. You may specify a single storage group, a list of groups, and/or a "group" of group names (e.g., STORGRP=DB*).

STORGRP= can be specified only for Fast COMPAKTion (TYPE=FASTCPK) or space release (TYPE=RLSE). You cannot specify STORGRP= and VOL= in the same run.

VOL= Specifies one or more disk volume serials which COMPAKTOR is to dynamically allocate and process. FROMDD= and VOL= may not both be specified. If VOL= is specified, there must be no DISK1 DD statement in the COMPAKTOR JCL.

For a COMPAKTion-from-backup (COMPAKT TYPE=COMPAKT or COMPAKT with no TYPE= operand), only one complete volume serial is allowed, e.g.,

VOL=TSO001

For TYPE=FASTCPK or TYPE=RLSE or SIMULATE, multiple volume serials may be specified:

- 1) A list of volume serials may be given, enclosed in parentheses, up to a maximum of 254, e.g.,

VOL=(TSO001,TSO002,TSO003)

COMPAKTOR will process the volumes in the order specified.

- 2) A volume group may be specified by placing an asterisk at the end of the volser prefix, e.g.,

VOL=TSO*

Within that group, COMPAKTOR will process the volumes in the order that their UCBs exist in the MVS system.

- 3) The two may be combined, e.g.,

VOL=(TSO*,PROD*,ABC001)

Each group will be processed in the order described above before going on to the next group or volume.

- 4) All online disk volumes may be selected by:

VOL=*

COMPAKTOR will process the volumes in the order that their UCBs exist in the MVS system.

COMPAKTOR will not check for duplicate volume serials, so volumes selected more than once will be processed more than once.

CONTINUED . . .

40.08 CONTINUED . . .

UNABLE=	<p>IGNORE – specifies that COMPAKTOR is to continue if some unusual condition occurs, such as very little free space on the volume, so that COMPAKTOR algorithms are unable to improve the volume.'</p> <p>By default, COMPAKTOR will issue a U0888 abend when it is unable to improve the volume.</p> <p>Innovation strongly recommends that you specify UNABLE=IGNORE.</p>
VSAMEXT=	<p>KEEP – specifies that all multi-extent ICF VSAM components are to retain the number and size of their extents. If used with VSRLSE=ALL, the unused space will be freed down to a CA (Control Area) boundary, but the extents occupied by the remaining used portion of the component will not be merged.</p> <p>MERGE – specifies that CPK is to attempt to merge most multi-extent VSAM components into a single extent.</p> <p>The default is MERGE except for a COMPAKT-from-backup with DUMP=NO, where the default is KEEP.</p>
VSRLSE= VSR	<p>Specifies that unused tracks within ICF VSAM components are to be freed. CPK uses the HI-USED-RBA field for the component to determine how much of the component is in use. If it is zero, CPK will not free any space from the data set.</p> <p>ALL – components of ICF VSAM data sets will be released to the CA (control area) boundary closest to the HI-USED-RBA; additional CAs may be retained as free space for expansion if %FREE= or %VSFREE= is specified. Space will not be released from ICF catalogs, the VVDS, or multi-volume components. See section 40.09 if space must be released from ICF catalogs.</p> <p>NO – space will not be released from ICF VSAM data sets, unless they are affected by the VSRLSE= operand of a SELECT statement.</p> <p>The default is NO.</p>
VTOC= V	<p>COMPAKT – specifies that the VTOC is to be COMPAKTed, embedded Format 0 DSCBs are to be eliminated. This option must be specified if you wish to change the location or size of the VTOC. VTOC=COMPAKT can be used only with a COMPAKT-from-backup (COMPAKT TYPE=CPK) and is invalid with TYPE=RLSE and TYPE=FASTCPK.</p> <p>NOCHANGE – specifies that the DSCB entries are to remain in the same location.</p> <p>Default is NOCHANGE. Innovation recommends VTOC=NOCHANGE unless you need to move or expand the size of the VTOC.</p>

CONTINUED . . .

40.08 CONTINUED . . .

%FREE= Specifies that for data sets from which CPK is releasing space, this percentage
%PSFREE= of the data set is to be left as free space after COMPAKTion. PSRLSE=,
%POFREE= PORLSE= and/or VSRLSE= must also be specified or else these parameters
%VSFREE= will be ignored.

CPK will only release space on PS, PO and ICF VSAM data sets and will never increase these data sets in size. For example, if a data set occupies 100 tracks of which 20 tracks are used, %FREE=50 will leave the data set with 40 tracks. Values may be from zero(0) specifying that all of the unused space is to be freed, to 99 which says that 99% of the data set is to be left free.

%FREE applies to PS, PO and ICF VSAM data sets. %PSFREE applies only to PS data sets, %POFREE only to PO data sets, and %VSFREE only to VSAM data sets. If %FREE is coded along with one or more of the others, it will apply to any of the data set types for which the specific keyword was not coded.

The default for all the keywords is 0.

40.09 SELECT STATEMENT

```

SELECT      DSN=filter          ,VSRLSE=NO|ALL
S           ALLDSN

           ,DSORG=(xx,xx...)    ,%FREE=nn

           ,EXTENTS=KEEP

           ,POS=VTOC|LVTOC|RVTOC|BEGIN|END|KEEP|ccccchhhh

           ,RLSE=NO|ALL|ROUND|TRK

```

SELECT STATEMENT This statement is used to define data sets requiring special treatment by COMPAKTOR. SELECT DSN= is also used to define data sets which are members of a sequenced set (see SEQUENCE statement).

NOTE: The SELECT statement specifies special processing options for selected data sets, but it does **not** imply that **only** those data sets will be moved or released by COMPAKTOR. Only those data sets that are identified as unmovable or active by COMPAKTOR plus data sets identified by a SELECT ...,POS=KEEP statement will be bypassed by COMPAKTOR.

The available options are:

Prevent multi-extent merges.	Control freeing of unused tracks.
Position the selected data sets.	

Up to 1000 SELECT statements may be input, but since each may select many data sets, an unlimited number of data sets may be affected.

OPERANDS **DSN=** Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). All data sets on the volume matching this name or filter will be affected by the options on this SELECT.

For ICF VSAM data sets, you must specify component names. CPK does not recognize cluster names.

EXAMPLES : DSN=USER1.JCL.CNTL
 DSN=**LIST
 DSN=PROD++.**.LIB*

NOTE: The DSG= operand documented in previous versions of COMPAKTOR is still accepted, but the DSN=operand with a generic data set name filter is the preferred way of selecting groups of data sets.

CONTINUED . . .

40.09 CONTINUED . . .

ALLDSN

Specifies that all of the data sets on the volume will be affected by the options on this SELECT.

Either DSN= or ALLDSN must be specified on a SELECT statement. If SELECT DSN= statements specifying a fully-qualified data set name are present in the control statement input stream, they must precede all SELECT DSN=filter or ALLDSN statements.

If a data set is made unmovable by the COMPAKTOR unmovable table or the list of currently active data sets (DSNENQ=), and a SELECT statement specifying RLSE= or POS= selects it, CPK will abend. If you do not wish to abend, code SELTERR=NO on the CPK control statement. If you wish the data set to be moved, code OVERRIDE=YES on the CPK control statement. [See Section 40.21](#) for special considerations on COMPAKTING active volumes.

DSORG=

Specifies that this control statement only applies to data sets with the specified DSORG(s).

PS – specifies physical sequential files (including extended format (striped) sequential).

PO – specifies partitioned files (including PDSE).

DA – specifies direct-access files.

EF – specifies VSAM files in an ICF catalog.

**EXTENTS=
E**

Specifies that all extents owned by the selected data sets are to be retained and not merged into a single extent, although they may be moved (use POS=KEEP to keep all the extents in their current locations).

If omitted, most multiple extent data sets are merged into a single extent, unless EXTENTS=KEEP is specified in the major statement.

40.09 CONTINUED . . .

POS=

If coded, specifies that positioning is requested for the selected data sets. Positioning is of two types: relative or absolute.

Relative positioning is specified by coding VTOC, LVTOC, RVTOC, BEGIN or END.

VTOC – specifies that the data sets are to be placed as near the VTOC as possible, on either side of it.

LVTOC – specifies that the data sets are to be placed as close to the VTOC as possible, on the left side (lower numbered tracks) of it.

RVTOC – specifies that the data sets are to be placed as close to the VTOC as possible, on the right side (higher numbered tracks) of it.

BEGIN – the data sets are placed as near the beginning of the disk volume as possible.

END – the data sets are placed as close to the end of the disk volume as possible.

Absolute positioning is specified by coding KEEP or an address (ccccchhhh).

KEEP – the data set is to retain its original position. If specified, RLSE=, VSRLSE= and %FREE = cannot also be specified.

ccccchhhh – code an 8- or 9-digit decimal cylinder and head address, when you wish to specify an absolute position for the first track of the data set. The first 4 or 5 digits are the decimal cylinder number (the first "c" may be omitted) and the last 4 are the decimal track number within that cylinder. Both cylinder and track addresses are relative to zero. POS=00000000 is invalid. CPK will fail if this address causes the data set to overlap with an unmovable data set or another absolute positioned data set. Absolute positioning is not valid when DSN=filter or ALLDSN are coded; a fully-qualified data set name is required.

If POS= is omitted, the position of the data set after COMPAKTion is determined by internal algorithms.

[See Section 40.12](#) for further considerations when positioning data sets. POS= must be omitted if the SELECT statement is part of a sequenced set, and it is invalid on TYPE=RLSE.

Note: Innovation does not recommend positioning data sets when using TYPE=FASTCPK.

CONTINUED . . .

40.09 CONTINUED . . .

RLSE=

Specifies the action to be taken for the unused tracks in the selected data sets. This option is ignored if the data set is not sequential or partitioned.

NO – unused tracks will be kept.

ALL – for track-allocated data sets, all unused tracks are to be freed. For cylinder-allocated data sets, only the unused cylinders are freed.

ROUND – only unused extents are to be freed.

TRK – all of the unused tracks are to be freed even if the data set is cylinder-allocated.

NOTE: If RLSE=TRK releases a cylinder-allocated data set past its cylinder boundary, CPK will reset the cylinder-allocated indication in the affected extent descriptors in the DSCB. However, the secondary allocation type will remain the same as before the release. (The secondary allocation type is shown in the SPACE ALLOC field in the CPK map). Some products, such as ISPF, test the secondary allocation type and may display a data set as occupying a number of whole cylinders when it really does not.

Default is the value of the PSRLSE= operand (if DSORG=PS) or PORLSE= operand (if DSORG=PO) of the SIMULATE or COMPAKT statement. If PSRLSE= and PORLSE= were omitted, then only data sets covered by a SELECT statement with the RLSE= operand will have free space released. If PSRLSE=/PORLSE= were specified, then SELECT statements with RLSE= can be used to override them for selected data sets.

VSRLSE=

Specifies the action to be taken for the unused tracks in any ICF VSAM components selected by this SELECT. It is ignored if the data set is not ICF VSAM.

NO – unused tracks will be kept.

ALL – components of ICF VSAM data sets will be released to the CA (control area) boundary closest to the HI-USED-RBA; additional CAs may be retained as free space for expansion if %FREE= or %VSFREE= is specified. Space will not be released from the VVDS or multi-volume components, but it will be released from ICF catalogs if DSN= specifies the fully-qualified catalog component name and OVERRIDE=YES was also specified.

The default is NO unless VSRLSE=ALL was specified on the major statement (except for ICF catalog components, which will be released ONLY if VSRLSE=ALL is specified on the SELECT).

%FREE=

Specifies that CPK is to leave this percentage of the data set free after COMPAKTion. RLSE= and/or VSRLSE= must also be specified or else %FREE= will be ignored. CPK will only release space on PS, PO and ICF VSAM data sets and will never increase these data sets in size. For example, if a data set occupies 100 tracks of which 20 tracks are used, %FREE=50 will leave the data set with 40 tracks. Values may be from zero (0) specifying that all of the unused space is to be freed, to 99 which says that 99% of the data set is to be left free.

Default is 0 unless %FREE= was coded on the CPK statement, or the %PSFREE=, %POFREE= and/or %VSFREE= operands were coded on the CPK statement, depending on the type of data set selected.

40.10 VTOC SELECT STATEMENT

```
SELECT      ***VTOC

            ,POS=KEEP|MIDDLE|BEGIN|END|ccccchhhh

            ,SIZE=nnnnn
```

VTOC STATEMENT This statement is similar in function and use to the SELECT statement, but applies only to the VTOC. You may use it to change the size and/or position of the VTOC. It may also appear as part of a sequenced set, under a SEQUENCE statement, in which case the POS= operand must be omitted; in this case the sequence set defines the data sets which will immediately precede and/or follow the VTOC. A SELECT ***VTOC can ONLY be specified with COMPAKT-from-backup (TYPE=CPK) and the VTOC=COMPAKT operand MUST also be specified.

OPERANDS *****VTOC** This is a required operand and identifies this SELECT statement as applying to the VTOC.

POS= Specifies the new position of the VTOC.

KEEP – the VTOC is to retain its position.

MIDDLE – you wish COMPAKTOR to position the VTOC as near the center of the disk volume as possible.

BEGIN – you desire the VTOC to be as near to the beginning of the disk volume as possible.

END – you desire the VTOC to be as near to the end of the disk volume as possible.

ccccchhhh – if you desire to position the VTOC at a specific location, code an 8 decimal number, the first 4 digits being the cylinder address, the last 4, the track address. Both addresses must be relative to zero. POS=00000000 is invalid. CPK will fail if the VTOC's new position overlaps an unmovable data set or another absolute positioned data set.

On a volume with over 65,535 tracks (such as the 3390-9) the **last track** of the VTOC must be no higher than relative track 65,535 (cylinder 4368, track 14 on a 3390-9).

If omitted, the default action is to retain the VTOC's position, unless this statement appears in a sequenced set (see SEQUENCE in [Section 40.11](#)).

SIZE= If coded, supply 1 to 5 decimal digits giving the VTOC size in tracks. A value of 0 is invalid, as is a value in excess of the device capacity. CPK will fail if the expanded VTOC overlaps an unmovable data set or an absolute positioned data set. If you reduce the size of the VTOC, make sure that it is large enough to hold all current active DSCBs. If not, COMPAKTOR terminates.

If omitted, the VTOC retains its original size.

CONTINUED . . .

40.10 CONTINUED . . .

**NOTES ON VTOC
POSITIONING**

When changing the location of the VTOC, you should consider the following:

On current-technology disks, with their fast access and cache, you may wish to place the VTOC at the beginning of the volume.

Place the VTOC in such a manner that its last track is also the last track of a cylinder, i.e., the VTOC ends on a cylinder boundary. This may also improve performance. When a VTOC is relatively positioned, as with POS=MIDDLE, POS=BEGIN, or POS=END, COMPAKTOR always places the VTOC to end on a cylinder boundary. This is also true when the VTOC is part of a sequenced set.

CONTINUED . . .

40.11 SEQUENCE AND ENDSEQ STATEMENT

SEQUENCE POS=VTOC|LVTOC|RVTOC|BEGIN|END|ccccchhhh
 SEQ
 (SELECT statements)
 ENDSEQ

SEQUENCE STATEMENT This statement is used when you wish to position a number of data sets, yet want to establish a positional relationship among the data sets. In essence, the SEQUENCE statement allows you to treat its member data sets as a single sequenced group for purposes of positioning. The SEQUENCE statement must be followed by one or more SELECT DSN= statements specifying fully-qualified data set names (no filters) to define the members of the set, and an ENDSEQ statement delimiting the end of the set. The positioning sequence for members of a set is implicit in the sequence of the SELECT statements; i.e., the first named data set is positioned first, followed by the second named data set, and so on. [See Section 40.12](#) for more details. A SELECT ***VTOC may also be included in the sequenced set to position the VTOC as part of the set, and to position data sets around the VTOC ([see Section 40.10](#) for information on moving the VTOC).

ENDSEQ STATEMENT This statement must be used to delimit a sequenced set. It must follow the last SELECT statement naming a member of the sequenced set.

**OPERANDS
 (SEQUENCE
 statement
 only)**

POS=

VTOC – the set is to be placed as near the VTOC as possible, on either side of it.

LVTOC – the set is to be placed as close to the VTOC as possible, on the left side (lower numbered tracks) of it.

RVTOC – the set is to be placed as close to the VTOC as possible, on the right side (higher numbered tracks) of it.

BEGIN – the set is to be placed as near the beginning of the disk volume as possible.

END – the set is to be placed as near the end of the disk volume as possible.

ccccchhhh – code an 8- or 9-digit decimal cylinder and head address, when you wish to specify an absolute position for the first track of the first data set in the set. The first 4 or 5 digits are the decimal cylinder number (the first "c" may be omitted) and the last 4 are the decimal track number which must always be 0000 since sequence sets always start on a cylinder boundary. The cylinder address is relative to zero. POS=00000000 is invalid. CPK will fail if this address causes any data set in the set to overlap with an unmovable data set or another absolute positioned data set.

The default is POS=VTOC.

NOTE: If the VTOC is itself part of the set (SELECT ***VTOC) POS=VTOC implies POS=MIDDLE: i.e. the set is placed as near the center of the disk volume as possible. If the VTOC is not part of the set, then POS=VTOC will place the set either before or after the VTOC; the VTOC will not be embedded in the set.

Including the VTOC as part of a sequence set implies that it will move, so this can be done only on a COMPAKT-from-backup (TYPE=CPK) with VTOC=COMPAKT specified.

WARNING: *Innovation does not recommend use of sequenced sets with FASTCPK since it may require more data movement and may increase the elapsed time of the COMPAKTion.*

CONTINUED . . .

40.12 COMPAKTOR DATA SET TYPES

COMPAKTOR recognizes the following data set types:

- TYPES**
- Unmovable data sets
 - Unchangeable data sets
 - Standard-user-label data sets
 - Cylinder-boundary data sets
 - Track-boundary data sets.

**UNMOVABLE
DATA SETS**

The following data sets are considered unmovable by COMPAKTOR:

- Any ISAM data set
- Non-ICF VSAM data sets
- ICF VSAM catalogs, page and swap data sets, and the VVDS if DUMP=YES and DSNENQ= are both specified, or if TYPE=FASTCPK is specified (see NOTE)
- All load libraries in the system linklist (LNKLSTxx or PROGxx) (see NOTE)
- Any data set marked unmovable in the DSORG field (i.e., DSORG=PSU, POU, or DAU) (see NOTE)
- Any BDAM data set whose Format 1 DSCB indicates that it is accessed via absolute addresses (OPTCD=A) (see NOTE)
- Any data set named in a SELECT statement also specifying POS=KEEP
- Any data set specified in the COMPAKTOR Unmovable Table (see NOTE)
- When using the data set enqueue option (DSNENQ=), any data set found to be active (see NOTE)
- If the CPKABSUN global option is set to YES ([see section 90](#)), any data set allocated with absolute track allocation (ABSTR) (see NOTE)
- Any SMS-managed sequential data set for which a checkpoint/restart is pending (see NOTE).

Data sets must be considered unmovable when they contain, within their data space, physical cylinder/track addresses. Moving these data sets would entail readjustment of these internal pointers, which COMPAKTOR is not designed to do. Note that a data set may also be unmovable if some other entity contains CCHH address pointers to locations within the data set. Unmovable data sets other than those which are inherently unmovable (such as ISAM) should be marked unmovable in the DSORG field (e.g., DSORG=PSU) for COMPAKTOR to recognize them.

On COMPAKTion-in-place processing, COMPAKTOR will not restore the data of unmovable data sets. Otherwise, COMPAKTOR handles unmovable data sets by always restoring them to the same locations, on the target disk volume, as they occupied on the dumped volume. In this manner, all physical pointers within or to these data sets retain their validity.

The following restrictions apply to unmovable data sets:

- You may not position most unmovable data sets. Data sets specified in the COMPAKTOR unmovable table or found active **can** be made movable using SELECT DSN= control statements if OVERRIDE=YES is specified on the CPK control statement.
- You may not free unused tracks in unmovable data sets.
- Unmovable data sets may not be members of a sequenced set.

NOTE: If a data set would be considered unmovable for one of the reasons that say “see NOTE”, you can force the data set to be moved by adding OVERRIDE=YES to the CPK control statement, and providing a SELECT statement which selects the data set.

CONTINUED . . .

40.12 CONTINUED . . .

**UNCHANGEABLE
DATA SETS**

The following data sets are considered unchangeable by COMPAKTOR:

- Any data set originally allocated using the ALX subparameter of the SPACE= keyword of a DD statement
- Any data set named in a SELECT statement specifying EXTENTS=KEEP
- All ICF VSAM clusters if VSAMEXT=KEEP is specified on a SIMULATE or COMPAKT major statement
- All non-VSAM data sets if EXTENTS=KEEP is specified on a SIMULATE or COMPAKT major statement

When COMPAKTOR recognizes a data set as being unchangeable, it never merges the data set's extents. However, it usually attempts to place all extents owned by the data set contiguous to one another.

**CYLINDER
BOUNDARY
DATA SETS**

These are data sets whose space was originally allocated so as to be on cylinder boundaries. The only restriction here is that these data sets must not be positioned on other than cylinder boundaries.

Although a data set was originally allocated on cylinder boundaries, it may later have been changed so that some of its extents are not on cylinder boundaries. When COMPAKTOR detects this, it treats such data sets as track boundary data sets, even though the DSCB specifies cylinder allocation.

**TRACK
BOUNDARY
DATA SETS**

These are data sets whose data space was originally allocated as either a number of tracks or a number of blocks. Absolute track allocated data sets are also considered track-boundary data sets. No restrictions apply to these data sets.

**DATA SET
ORGANIZATIONS**

Other than the restrictions just described (unmovable data sets, etc.), COMPAKTOR can move all types of data sets, including:

- PS (Physical Sequential)
- PSE (Physical Sequential - Extended Format (Striped and compressed SMS data sets))
- PO (Partitioned)
- POE (PDSE and HFS data sets)
- DA (Direct Access - BDAM)
- ICF VSAM (including IMS and DB2 data bases)

COMPAKTOR can release space from all of the above except DA.

40.13 COMPAKTOR DATA SET POSITIONING

This section details some things to keep in mind when using the COMPAKTOR data set positioning facility. This facility allows you to position single data sets, groups of data sets and sequenced sets.

NOTE: Innovation does not recommend explicit data set positioning in most circumstances.

Explicit positioning limits the ability of COMPAKTOR to do its job. Current disk technology, with faster access and caching, makes most data set positioning unnecessary. For further information on COMPAKTOR's data set positioning facility, see the member \$\$CPK in the FDR Installation Control Library (ICL).

CONTINUED . . .

40.14 COMPAKTOR REORGANIZATION

COMPAKTOR uses a number of algorithms in order to determine the optimal method of reorganizing a disk volume. CPK attempts to produce a reorganized volume with as few free space areas and as few multi-extent data sets as possible. These algorithms are not used for TYPE=RELEASE since no data sets are moved.

FREE SPACE AREAS A 'free space area' is an unused area of a volume, unallocated cylinders or tracks. All COMPAKTOR algorithms attempt to minimize the number of free space areas in a reorganized volume (see SIZEKEEP= in this section).

MULTI-EXTENT DATA SETS COMPAKTOR attempts to reduce most multi-extent data sets to a single extent (see SIZEKEEP= in this section). When this is not possible, it leaves data sets into 2 or more extents. However, no data set is ever split into more extents than it originally possessed; i.e. no data set becomes 'worse' than before COMPAKTion. In extreme cases, CPK may change the size of the extents. EXTENTS=KEEP and/or VSAMEXT=KEEP can override this.

COMPAKTOR ALGORITHMS Before any algorithm is used, COMPAKTOR always does the following:

- placement of unmovable data sets
- placement of VTOC, if unmovable or absolute positioned
- placement of absolute positioned data sets and sequenced sets
- placement of relative positioned VTOCs and sequenced sets
- placement of relative positioned data sets.

COMPAKTOR then goes through various algorithms to determine the placement of the remaining data sets, finally choosing the algorithm which results in the fewest free space areas.

EFFECT OF THE SIZEKEEP OPTION The SIZEKEEP= performance option may affect the results of the COMPAKTOR algorithms just described. The SIZEKEEP= operand has a default (specified in the FDR Global Options Table) which may affect results even if the operand is not present. [See Section 40.08](#) for the further description of SIZEKEEP=.

The SIZEKEEP= operand has 3 subparameters, a size (in tracks), a percentage and an extent count, which default to 100 tracks, 90% and 60 extents (unless you have changed the defaults in the FDR Global Options Table). The purpose of SIZEKEEP= is to improve the efficiency of COMPAKTOR by moving/restoring as little data as possible, as long as a significant improvement can still be made in the fragmentation of the volume. The percentage is a minimum reduction in the number of free space areas on the volume which must be produced for SIZEKEEP= to be honored; so by default, COMPAKTOR must be able to reduce the free area count by 90%.

The size subparameter of SIZEKEEP= causes COMPAKTOR to look for contiguous groups of allocated tracks which are at least that large (by default, 100 tracks or larger). These groups may contain multiple data sets, but they are bounded by an existing free space area on either side. COMPAKTOR will mark these groups unmovable and then execute each placement algorithm to move other data sets and see if the required percentage reduction in free areas can be achieved. If so, an algorithm will be chosen. If not, then part of that list of unmovable groups is made movable (the smaller groups) and the algorithms are driven again. This repeats until the required percentage is achieved or until the list is empty.

The third subparameter, extent count, controls SIZEKEEP= processing of multi-extent data sets. Data sets with more than the specified number of extents (default 60) will not be considered to be part of any SIZEKEEP track group, so they will be movable. If the unmovable groups chosen by SIZEKEEP= contain one or more extents of a multi-extent data set, but not all of its extents, COMPAKTOR may move the other extents of the data set but will not combine them. In other words, the actions of SIZEKEEP= may cause COMPAKTOR to leave some data sets in multiple extents where it would normally combine them.

CONTINUED . . .

40.14 CONTINUED . . .

**EFFECT OF THE
SIZEKEEP
OPTION
(continued)** You can specify SIZEKEEP=0 to nullify the operation of SIZEKEEP, but this may significantly increase the elapsed time of the COMPAKTion while providing very little additional benefit. For Fast COMPAKTion (TYPE=FASTCPK) SIZEKEEP=0 may still keep some groups over 100 tracks as long as the free space areas are reduced to 1 or 2.

Innovation recommends that in most cases you use the default SIZEKEEP of 100,90,60.

40.15 COMPAKTOR OUTPUT

When a disk volume is COMPAKTeD, COMPAKTOR prints 'before' and 'after' maps of the disk volume on the SYSMAP DD statement. This is true whether you actually COMPAKTeD the volume or only simulated COMPAKTion. Along with these detailed maps, which can be suppressed by the MAPS= operand, a summary page for each volume is also printed. At the end of the COMPAKTOR run, a combined summary is printed (on DD SYSSUMM or SYSMAP) giving a 2 line summary for each disk processed, sorted by volser. [See section 40.22](#) for examples of these reports.

**MEASURING
THE
IMPROVEMENT**

To determine how much COMPAKTOR has improved a volume, you can review the summary or the combined summary and compare the before/after figures for the following fields:

NUMBER OF FREE SPACE AREAS. This is a measure of volume fragmentation and will usually be smaller after COMPAKTion (will be larger after RELEASE).

NUMBER OF FREE CYLINDERS. This too is a measure of volume fragmentation. This value will usually be larger after COMPAKTion.

SIZE OF LARGEST FREE AREA, IN TRACKS. Again, a measure of volume fragmentation. Normally, you will find a larger contiguous free area after COMPAKTion.

NUMBER OF DATA SETS WITH 2 OR MORE EXTENTS. This figure tells you how many multi-extent data sets exist on the volume. Will almost always be less after COMPAKTion.

NUMBER OF EMPTY TRACKS IN PS/PO/VSAM DATA SETS. This figure is the total number of unused but allocated tracks on a volume. This value will be smaller if you have specified operands to release unused space from these data sets.

IBM FRAGMENTATION INDEX. This index attempts to show the level of fragmentation on a disk volume, calculated by an IBM-provided formula using the size of free space areas on the volume. A value of 0.000 is achievable only when there is only a single free area, and values close to 1.000 occur when there are many very small free areas (such as every other track). The formula is a non-linear function involving natural logarithms and is biased toward most of the free space being collected into one large area. Innovation believes it is difficult to understand and does not always accurately reflect the effect of fragmentation on the ability to allocate data sets; generally, if there are 3 or fewer free areas the fragmentation index should be ignored. Innovation suggests that volumes with indexes over 0.200 or over 20 free space areas should be COMPAKTeD, but you may need to run COMPAKTOR simulations on all your volumes to decide what threshold should be used by your installation.

**SYSPRINT
MESSAGES**

Several messages in the SYSPRINT DD output indicate the amount of work that COMPAKTOR did on each volume:

- CPK316I shows how many tracks COMPAKTOR actually restored, moved, or released
- CPK319I shows how many tracks COMPAKTOR was able to treat as unmovable due to SIZEKEEP= options, allowing you to gauge the effectiveness of SIZEKEEP= (the larger the value, the less work CPK did). If the message is absent, either SIZEKEEP=0 was specified or COMPAKTOR was unable to meet its objective without making all tracks movable.

**UNABLE TO
IMPROVE**

There are several conditions under which COMPAKTOR may not be able to improve the condition of a volume:

- User-specified positioning forced COMPAKTOR to create a more fragmented volume
- An excessive number of unmovable data sets may prevent improvement
- The volume is already optimally organized. Perhaps you just COMPAKTeD it or perhaps there is only one data set on the volume.

If any user-specified data set or VTOC positioning was performed or if any unused tracks were freed, COMPAKTOR will go ahead and COMPAKT the volume. If not, message CPK556W UNABLE TO IMPROVE is issued and the volume is bypassed (to force COMPAKTion, rerun specifying space release or positioning parameters).

CONTINUED . . .

40.16 COMPAKTOR TECHNICAL DESCRIPTION

This chapter deals with a few items of a technical nature which should be considered when using COMPAKTOR.

**ACCESS
TECHNIQUE**

COMPAKTOR uses EXCP level code to access disk devices. COMPAKTOR always reads or writes one or more entire tracks at a time. Hence, it is access-method independent. Individual data sets are not opened.

If you are also licensed for FDR InstantBackup and execute FASTCPK on a EMC Symmetrix disk, you can add the EMCCOPY=YES operand to the COMPAKT statement to cause COMPAKTOR to use an internal EMC copy feature to move tracks without sending them to and from the CPU, resulting in a reduction in elapsed time. [See Section 25](#) for more information on FDR InstantBackup with EMC disks.

If you are also licensed for FDR InstantBackup and execute FASTCPK on a IBM RVA or StorageTek Iceberg or SVA disk with the Snapshot feature, you can add the SNAPSHOT=YES operand to the COMPAKT statement to cause COMPAKTOR to use the Snapshot feature to quickly move tracks without sending them to and from the CPU. [See Section 40.26](#) for more information on FASTCPK with the RVA/Iceberg/SVA. [See Section 26](#) for more information on FDR InstantBackup with RVA/Iceberg/SVA.

**VOLUME
LABEL
PROCESSING**

As a default, COMPAKTion from an FDR backup restores the volume serial number of the dumped volume. If either the VTOC location or the volume serial number of the output volume is changed, COMPAKTOR automatically updates the UCB to reflect the changes. If another mounted volume has the same volume serial number COMPAKTOR will place this volume offline. In that case, you may want to vary the original volume offline and then mount the new volume in its place. Optionally, you may request that COMPAKTOR retain the volume serial number of the target volume; however, if the volume contains ICF VSAM clusters or is SMS-managed, this may make the volume unusable.

**VOLUME
TYPES**

COMPAKTOR only supports validly formatted OS direct-access volumes. VSE volumes with VTOCs starting at cylinder zero head zero and VM/CMS formatted disk volumes are not supported. COMPAKTOR automatically supports volumes with a non-standard size (a number of cylinders not matching the size of a standard IBM DISK) by honoring the number of cylinders shown in the VTOC of the output disk. This supports OS-formatted VM mini-disk and non-IBM disk subsystems with a non-standard size.

DIAGNOSTICS

COMPAKTOR performs a unique VTOC validation function. All VTOC errors and/or inconsistencies detected cause appropriate diagnostic messages to be issued.

Upon termination of a COMPAKTOR job step, either a completion code of zero is provided or one of a number of user abends is issued. User abends are used rather than non-zero completion codes to ensure that attention is drawn to all diagnostic messages. Non-zero completion codes often tend to be overlooked.

User abend U0888 is used to indicate that COMPAKTOR encountered an unusual condition, but that it did not stop COMPAKTOR from completing the COMPAKTion. The abend is issued to call attention to the message documenting the error, since it may have affected the usability of the COMPAKTed volume. However, the option CPKCC in the FDR option table may be changed to substitute a non-zero completion code for the U0888 abend .

CONTINUED . . .

40.16 CONTINUED . . .

**UPDATE
INDICATORS**

After a volume has been COMPAKTed, the update indicator will be turned on in the Format 1 DSCB of every data set that COMPAKTOR moved on the volume, as if those data sets had been opened for output.

This is done for compatibility with ABR incremental backup. Because of the "backward recovery" technique used by ABR full-volume recovery ([see Section 50](#)), a data set must be backed up by ABR if its location changes, even if its contents have not changed. This means that the next ABR incremental backup (DUMP TYPE=ABR) done on a volume after a COMPAKTion will backup many of the data sets on the volume.

NOTE: if you are using ABR incremental backup and also use IBM's DFDSS DEFRAg rather than FASTCPK, contact Innovation for special consideration.

The SIZEKEEP= option of COMPAKTOR is designed to reduce the impact of this. SIZEKEEP attempts to avoid moving the largest data sets on the volume as long as it can still achieve a significant reduction in the number of free-space areas on the volume by moving the smaller data sets. If it is successful, then only the smaller data sets will have their update flag set and the amount of data to be backed up by the ABR incremental will be significantly less.

In any case, Innovation recommends running FASTCPK before the ABR full-volume backups to avoid the problem altogether.

NOTE: A Custom Zap not to turn on the update bit is available for CPK customers who do not have ABR.

SECURITY

Before modifying a volume, COMPAKTOR checks to see if any other jobs in the system have OPEN data sets on the device to be restored by testing a OPEN count in the UCB of the disk device. If so, COMPAKTOR may issue a WTOR for console message FDRW81 and prevent modification until all other users of the device have closed their open data sets. However, the operator has the option to allow the COMPAKTion to proceed despite the active data sets. This is not desirable unless the rules for COMPAKTing an active volume have been followed ([See Section 40.21](#)). This WTOR can be suppressed by the option ACTMESS=NO on the COMPAKT statement or in the FDR/ABR option table. Fast COMPAKTion (TYPE=FASTCPK) and space release (TYPE=RLSE) default to COMPAKTing active volumes and assumes ACTMESS=NO.

A WTOR for console message FDRW80 is normally issued to allow the system console operator to confirm a CPK operation; if the COMPAKTion is undesirable, the operator may cancel it at this point. This WTOR may be suppressed by the option CONFMESS=NO on the COMPAKT statement. Fast COMPAKTion and space release assume CONFMESS=NO.

RACF volume protection (CLASS= 'DASDVOL') is supported by COMPAKTOR if the ALLCALL option is enabled in the FDR option table. Individual data sets are not checked.

COMPAKTOR supports the FDR OPEN EXIT. Control is passed to the user exit prior to any modification of the volume. The COMPAKTion or release can be terminated by this exit.

40.16 CONTINUED . . .

**SHARED DASD
ENVIRONMENT**

When COMPAKTING a disk volume in a shared DASD environment, you must be extremely careful. Unless you have a cross-system ENQ facility (GRS or MIM from Computer Associates), COMPAKTOR has no way to know what data sets are in use on another CPU or to prevent other CPUs from OPENing data sets in the middle of COMPAKTion. In this case, you may need to manually ensure that the volume is OFFLINE to the other systems in the shared DASD environment.

Note that the default of ENQ=RESERVE on the COMPAKT statement does not fully protect you. Assume you have two CPUs, System A and System B, which share a number of DASD devices. Also, assume that a long-running job is executing in System B, accessing a data set residing on a volume mounted on one of the shared devices. If you now start a job in System A in order to COMPAKT the same volume and you specify ENQ=RESERVE, the following occurs:

1. COMPAKTOR takes ownership of the device via the RESERVE macro, forcing the job in System B to wait until the device is again available.
2. The volume on the device is COMPAKTeD, and the data set which was being accessed from System B is moved to a new disk location.
3. And now, after COMPAKTOR issues a DEQ macro to release the device, the job in System B, unaware that its data set has been moved, tries to access it at its old location. Results are unpredictable and certainly undesirable.

If you have a cross-system ENQ facility, and that facility is broadcasting SYSDSN ENQs to other CPUs, then you can use the DSNENQ= option (as described in [Section 40.21](#) for COMPAKTING active volumes) to identify the data sets in use on the other systems and make them unmovable. If that facility also broadcasts SYSVTOC ENQs, it may be possible to have that facility convert the hardware reserve done by CPK into a global ENQ instead, so that only the VTOC will be ENQed rather than the entire volume being RESERVEd; however, the vendor documentation for your cross-system ENQ facility should be consulted before allowing SYSVTOC to be converted as there may be additional considerations.

NOTE: additional information is available on cross system ENQ considerations, in the FDR Installation Control Library (ICL) under the member name ENQ.

40.17 COMPAKTOR VTOC ERRORS

As previously noted, COMPAKTOR performs rather stringent error diagnostics on VTOCs. Problems that exist but are not diagnosed by any other software are detected and reported by COMPAKTOR. Our experience indicates that these are the most common errors:

- INVALID F5/F7 DSCBs** Format 5 and 7 DSCBs describe the free space on a volume with no indexed VTOC (F7s are used on large volumes such as 3390-9s). There generally are 2 types of F5/F7 DSCB errors: either the free space described is not valid (starting or ending track address is not valid), or a free area described is actually allocated to a current data set. The first case leads to loss of free space, but is generally non-harmful; the second case, however, is extremely serious and can lead to loss of data if the space is assigned to a second data set.
- NOTE:** If an active indexed VTOC exists on a volume, then the F5/F7 DSCBs are never valid, and CPK does not inspect them.
- OVERLAPPING EXTENTS** Overlapping extents occur when two or more F1 or F3 DSCBs define extents which contain the same tracks. This is an extremely serious condition and can lead to loss of data.
- INVALID EXTENT DESCRIPTORS** An extent descriptor in an F1 or F3 DSCB is invalid if its cylinder and track addresses are invalid for the device or the extent ending address is lower than the starting address. Such a data set may be unusable.
- INVALID FLAGS IN EXTENT DESCRIPTORS** Each extent descriptor has a flag byte, describing the type of extent. A common error is that the flag byte defines the extent as being on a cylinder boundary when it is not. This error may or may not be serious, depending on data set organization and access techniques used.
- BROKEN F1-F2-F3 DSCBs** Format 2 DSCBs exist for ISAM data sets; Format 3 DSCBs exist for any type of data set which has more than 3 extents. The Format 1 DSCB points to the Format 2 or Format 3 DSCBs if they exist (A Format 2 may point to a Format 3). COMPAKTOR detects any broken chains, as well as any unrelated F2/F3 DSCB; i.e. when an F2/F3 DSCB is present but its F1 DSCB is missing. This condition is serious if a broken chain exists, but harmless if an unrelated DSCB exists. However, the error must be corrected for COMPAKTOR to run.
- See member \$\$CPK in the FDR ICL (Installation Control Library) for information on correcting DSCB chains.
- FORMAT 4 DSCB ERRORS** A warning message is issued if error flags are found set in the VTOC status byte field of a Format 4 DSCB. Also, if the F4 DSCB is not the first record in a VTOC, COMPAKTOR cannot process the volume; the volume violates IBM Operating Systems standards.

CONTINUED . . .

40.18 COMPAKTOR ERROR RECOVERY

COMPAKTOR has been designed to handle various error conditions by recovering from the errors. Recovery actions are as follows:

VTOC ERRORS COMPAKTOR recovers from certain VTOC error conditions, such as invalid F5 DSCBs and invalid flags in extent descriptors, automatically.

I/O ERRORS I/O errors while reading/writing a VTOC or the first few blocks of an FDR dump tape (the DSF blocks) cause immediate abnormal termination. This is because the involved volumes are essentially unusable.

DISK I/O ERRORS When operating system error recovery notifies COMPAKTOR of a permanent I/O error while writing a disk track during the restore, COMPAKTOR skips that track. Up to 20 tracks in error may be skipped. At the end of the restore operation, COMPAKTOR prints a list of all non-restored, or skipped tracks; the data sets involved can be determined by comparing this list to the "before" map. It is then up to you to determine whether the restored volume is usable or not.

A FASTCPK will terminate when even one disk I/O error occurs. If the error is in the early stages of FASTCPK, no harm has been done and the COMPAKTion will simply terminate. If changes have been made in the volume, COMPAKTOR will automatically enter recovery (described below) and return the volume to its original state.

TAPE I/O ERRORS When COMPAKTion-from-backup is notified of a permanent I/O error while reading from the backup data set by system error recovery, it skips the bad tape blocks and attempts to restore as much of the volume as possible. Each tape block may contain one or more original disk tracks. A list of non-restored tracks is printed after the end of the restore operation; the data sets involved can be determined by comparing this list to the "before" map. It is then up to you to determine whether the restored volume is usable or not. Up to 20 tape blocks may be skipped.

USABLE VOLUMES When mapping or simulating, all volumes are always usable, since COMPAKTOR never alters them.

When COMPAKTing, a volume is usable if:

1. Message CPK321I was not issued prior to termination.
2. If message CPK322I was issued prior to termination. However, if termination was abnormal, the volume may or may not be usable.

If LOG=YES was specified, CPK also issues message FDRW82 to the operator concurrent with the CPK321I and CPK322I messages, so that the operator knows when CPK has begun modifying the volume and when it is done.

I/O ERROR DIAGNOSTICS When a tape or disk I/O error is encountered, COMPAKTOR prints a diagnostic DUMP to aid you and us in identifying the track or block in error and the cause of the error.

FASTCPK RECOVERY Fast COMPAKTion (FASTCPK) maintains a log on the disk in which all track movement is recorded. If any error condition prevents FASTCPK from completing, it automatically enters a recovery routine which moves every track back to its original location, returning the volume to its original state. Note that this recovery routine is not as efficient as the actual COMPAKTion and may take several times as long as it took COMPAKTOR to get to the point of the error.

If the FASTCPK job is prematurely terminated with no opportunity to recover (such as a system crash), another FASTCPK job run against the volume will recognize that this occurred and will use the recovery log to perform the same recovery, automatically. A CPK MAP job will notify you that the recovery is pending.

CONTINUED . . .

40.19 COMPAKTOR PROBLEM DETERMINATION

Although COMPAKTOR has been engineered to have a high degree of reliability, sometimes errors can occur which necessitate problem determination and correction.

VTOC ERRORS When VTOC errors are diagnosed by COMPAKTOR, we suggest you do the following:

- Determine the true error cause. COMPAKTOR will usually inform you of the address of the DSCB in error. You can then display the VTOC using either the IEHLIST utility (statement LISTVTOC with the DUMP option), or FDRDSF with the PRINT TYPE=DSF,DSN=VTOC option ([See Section 20](#)).
- If the error can be fixed by COMPAKTOR, we suggest you COMPAKT the volume.
- If COMPAKTOR cannot correct the error, you must fix the invalid DSCB by using the IBM SUPERZAP utility program as described in IBM publications under the name AMASPZAP. Member F3CLEAR in the FDR ICL (Installation Control Library) is an example of running AMASPZAP to correct a VTOC.
- If you are uncertain of the nature of the error or if you just feel insecure, do not hesitate to contact Innovation Technical Support.

I/O ERRORS When I/O errors occur, we suggest you carefully examine all COMPAKTOR output and attempt to determine the cause of the error. Areas of importance are:

- The IOB sense information
- The CSW in the IOB
- The last disk track accessed – track in error
- The CCW being executed.

CONTINUED . . .

40.20 RECOVERING FROM COMPAKTOR ERRORS**RESTARTING
COMPAKT
FROM BACKUP**

The following is a description of the condition of the disk volume if COMPAKTOR is unable to complete a DUMP=YES COMPAKTion due to system crash, operator cancel or any other abnormal termination. (A DUMP=NO COMPAKT-from-backup may simply be restarted).

First COMPAKTOR will dump the volume to the backup file using FDR. If a termination occurs during this phase, the volume has not been modified in any way. The COMPAKTion can be restarted. After a successful completion of the backup, COMPAKTOR will map the volume to document current data set placement (the "before" map). If a termination occurs during this map, you can also restart the COMPAKTion since the volume has not been modified.

If COMPAKTOR determines that the volume can be improved by COMPAKTion, an operator message (FDRW80) is issued (unless the CONFMESS=NO option has been selected). If the operator responds YES to this message, COMPAKTOR will print CPK321I COMPAKTOR RESTORE STARTED. From this point until message CPK322I COMPAKTOR RESTORE ENDED is issued, the volume is being modified and any failure may leave the volume in a corrupted state. If LOG=YES was specified, the operator will be notified by an FDRW82 Message when the restore starts and ends.

First, COMPAKTOR resets all of the extent information in the VTOC to reflect the new locations of all of the data sets. In the Format 4 DSCB, the key field may be set to binary zero, the pointer to the highest address of the FORMAT 1 DSCB (high water mark) may be set to the second DSCB in the VTOC, and the count of available FORMAT 0 DSCBs may be set to zero. The Format 5 may be set to indicate that there is no free space, and the indexed VTOC, if any, is disabled.

The data tracks are now restored from the backup to their new locations on the disk. The data is in cylinder and head sequence on the backup according to the old locations of the data sets.

After the data is restored, COMPAKTOR will reset the Format 4, rebuild the free space pointers and issue the restore completed message (CPK322I). The volume is now usable again.

Finally, the VTOC index is rebuilt (if necessary) and the UCB is updated.

If the volume is accessed after a failure during the restore phase, between the CPK321I and CPK322I messages, certain functions may seem normal. Any data sets which were unmovable or completely restored will be accessible. However, all others may list normally in the VTOC but their data has not been relocated. There is no easy way to tell which data sets have been restored and which have not. Users may not be able to allocate new data sets on the volume. In addition, the volume may not be mountable during IPL, or on a shared system.

COMPAKT-from-backup with DUMP=YES should never be restarted after a failure in the restore phase. The safest procedure is to do an FDR full volume restore from the backup created during the COMPAKTion.

Normally an FDR restore can be executed after a failed COMPAKTion during the restore phase. If FDR abends with an Sx13 on the VTOC, a re-initialization of the volume is required, prior to the restore.

If you were COMPAKTing an active volume (that is, if the operator replied IGNORE to message FDRW81 or if the COMPAKT statement specified ACTMESS=NO) and there were data sets being updated during COMPAKTion, recovery is more difficult. If a full volume restore is not desirable, you can restore any data set which was relocated, from the backup tape. A DSF restore should be used specifying EXCLUDE for all of the data sets which may have been updated since the backup (data sets may have been updated if they were listed in FDR158 messages, or are included in the COMPAKTOR unmovable table) followed by a SELECT ALLDSN. To reset the VTOC you set the DOS flag x '80' in the DS4VTOCI field (disp. 58(X '3A')) using SUPERZAP. The operating system will recalculate the free space when any new data set is allocated to the volume. The VTOC index, if present, must be rebuilt by using ICKDSF.

CONTINUED . . .

40.20 CONTINUED . . .

**RESTARTING
COMPAKT
FROM BACKUP
(continued)** If COMPAKTOR is restarted after a failure during the restore phase, with DUMP=YES, you risk overwriting the very backup you need for the restore. The dump will take place and usually COMPAKTOR will issue the message CPK556W UNABLE TO IMPROVE. If the operator mounts the same volumes (Tape Management Systems may allow the rewrite with the same data set name), the new backup will be written over these volumes.

A duplicate tape copy (TAPE11) is recommended in the event the first copy is damaged during the restore phase. The restore can then be done from the duplicate copy.

**RESTARTING
FAST
COMPAKTION** A Fast COMPAKTion (TYPE=FASTCPK) goes through many of the same steps described earlier in this section for COMPAKT-from-backup. However, Fast COMPAKTion simply moves tracks on the volume until all tracks have reached their target locations. During this process, a failure may leave the volume in a corrupted state since the VTOC may not point to the current locations of the data for many data sets.

Fast COMPAKTion includes an automatic recovery function. At the point that movement of data actually begins, a recovery data set is allocated on the volume with the name:

COMPAKT.IN.PROGRESS.OR.RECOVERY.Dyyddd.Thhmm

which is a few tracks in length (depending on the number of tracks being moved). "yyddd" is the Julian date the data movement started and "hhmm" is the starting time. If you see this data set in the VTOC of a volume, and there is no Fast COMPAKTion running against that volume, then a Fast COMPAKTion failed and the volume is in recovery status. If you run any COMPAKTOR function against a volume in recovery status (even a MAP) you will receive a CPK586W message warning you that recovery is required (only Fast COMPAKTion will be allowed to complete).

As it proceeds, Fast COMPAKTion records restart information in that recovery data set. If the Fast COMPAKTion is interrupted, you simply need to rerun a Fast COMPAKTion against that volume; COMPAKTOR will recognize that it was interrupted and automatically restarts. Usually it will recover by reversing the track movement to return all tracks to their original locations as if the COMPAKTion had never been done. If it failed in the final stages of its operation, it will attempt to complete the data movement.

You must restart the Fast COMPAKTion as soon as possible after the interruption.

Except for data sets that were active at the start of the original Fast COMPAKTion and other data sets that were not moved, the volume will not be usable until COMPAKTOR is allowed to complete. To ensure that no new data sets are allocated or old data sets extended (which might make recovery impossible), COMPAKTOR will change the VTOC and VTOC INDEX (if present) to indicate that there are no free tracks on the volume, until the COMPAKTion successfully completes. You must not attempt any corrective operations of your own on the volume, since this may result in permanent loss of data.

If I/O errors occur during Fast COMPAKTion, COMPAKTOR will immediately go into recovery mode and attempt to return the volume to its original configuration.

**RESTARTING
SPACE
RELEASE** Since COMPAKTOR space release (TYPE=RLSE) does not move any data, but simply updates the VTOC and VVDS to indicate released space, the release job can simply be restarted.

CONTINUED . . .

40.20 CONTINUED . . .

**CANCEL
PROTECTION**

Since the majority of interrupted COMPAKTOR jobs result from operators issuing CANCEL commands, COMPAKTOR includes a CANCEL protection function. If a CANCEL is issued against a COMPAKTOR job at any time during the restore phase (while the VTOC, VVDS, or data tracks are actually being updated, moved, or restored), COMPAKTOR will intercept that CANCEL and issue an FDRW99 Message to the operator indicating that the CANCEL will result in volume corruption. The operator has the option of ignoring the CANCEL, cancelling the job, or bringing COMPAKTOR to a graceful termination when it finishes processing the current volume. The operator should select the option to cancel the job only if COMPAKTOR appears to be in a loop.

With CANCEL eliminated as a cause of interruptions, only hardware and software system failures remain as potential sources of interruptions, and these are increasingly rare.

40.21 COMPAKTING ACTIVE VOLUMES

Although COMPAKTOR enqueues on the VTOC of the target volume and checks the use count for the device in the UCB, some system data sets are opened when the Operating System is initialized; they are not all accounted for by the UCB use count.

ACTIVE DATA SETS

The following data sets are a partial list of active data sets which may not be enqueued by the operating system so that COMPAKTOR is unable to tell if they are in use. If applicable, these data sets and any other data sets which are not enqueued should be put in the unmovable table. Failure to do so may result in system failures.

Add These data sets to the COMPAKTOR unmovable table

JES PROCEDURE LIBRARIES (PROCLIBS)
 JES SPOOL and CHECKPOINT DATA SETS
 TAPE MANAGEMENT SYSTEM DATA SETS
 CICS JOURNALS
 SYS1.BROADCAST
 SYS1.MANx (SMF) DATA SETS
 OS CVOL CATALOGS (SYSCTLG)
 SECURITY SYSTEM DATASETS
 FDR/CPK PROGRAM LIBRARY

DATA SET ENQUEUE

When COMPAKT-from-backup is executed with DUMP=YES, FDR will determine which data sets are active if the data set enqueue option (DSNENQ=) is specified on the DUMP TYPE=FDR statement. A data set is considered active if it is enqueued with a QNAME of 'SYSDSN' and an RNAME of the data set name. This is the enqueue used by the job scheduler. In addition, ICF VSAM catalogs, page and swap data sets, and the VVDS are always considered active with this option.

For Fast COMPAKTion (TYPE=FASTCPK) and space release (TYPE=RLSE) the DSNENQ= operand is supported on the COMPAKT statement. However, it defaults to DSNENQ=USE so these operations automatically recognize active data sets.

However, there is one circumstance in which you do NOT want data set ENQs. The ENQs done by the operating system are strictly on the data set name; it is not possible to tell on which volume the data set in use resides. If you are COMPAKTING a volume which is a duplicate of another active volume (such as when you are creating a duplicate SYSRES volume), DSNENQ= would consider any data sets which were in use on the original volume to be unmovable on the duplicate volume, even though they are not really in use. If DSNENQ=NONE is specified, COMPAKTOR will be able to do a better job of COMPAKTING the duplicate volume.

UNMOVABLE DATA SETS

COMPAKTOR supports a table of unmovable data sets. SYS1.VTOCIX, SYS1.VVDS and SYS1.LOGREC are automatically included in this table. You may specify additional data set names or groups to be considered unmovable. Entries can be added to this table using the FDR ISPF support (Section 92) or by executing program FDRZAPOP. Data sets specified by the unmovable table or POS=KEEP or found active by the data set enqueue option will not be moved or restored by COMPAKTion using the DUMP=YES option or by FASTCPK or RELEASE.

NOTE: OVERRIDE=YES enables you to move data sets in the unmovable table or active data sets.

SYSRES

COMPAKTOR will not allow you to modify the active system residence volume of the system on which CPK is running.

COMPAKTING ACTIVE VOLUMES

If you must COMPAKT a volume which has data sets in constant use, you can do so. Fast COMPAKTion (TYPE=FASTCPK) is designed to process active volumes by default. Do not override the default of DSNENQ=USE. This also applies to space release (TYPE=RLSE). However, the improvement that CPK can achieve on an active volume is limited because many data sets are made unmovable.

WARNING: Do not use COMPAKT-from-a previously created backup (TYPE=COMPAKT) to COMPAKT active volumes.

CONTINUED . . .

40.21 CONTINUED . . .

OVERRIDING THE UNMOVABLE TABLE OR ACTIVE LIST Normally COMPAKTOR will abend if a SELECT DSN= control statement was specified for a data set found in the unmovable table or found to be active by the data set enqueue option. If OVERRIDE=YES is specified on the CPK control statement SELECT control statements will be honored for these data sets unless the data set is truly unmovable. This should be done with care. For example: ICF VSAM catalogs may remain OPEN for a considerable period of time. If they are moved while OPEN, the results are unpredictable. If an active tape management or linklist data set is moved, the system will probably crash.

WARNING: If the SYS1.VVDS, PAGE, SWAP, VTOC, Indexed VTOC or ICF catalogs are moved, you must vary the volume offline and remount the volume after COMPAKTion on all the CPUs which can reference this volume including the one CPK was run on. If this is not done and any of these data sets were previously opened by the system, VSAM will continue to use the old location. The DSNENQ= option will cause COMPAKTOR to consider page,swap,and ICF catalogs to be unmovable.

SIMULATION OF AN ACTIVE VOLUME If a SIMULATE is executed against an active volume, FASTCPK and RELEASE will NOT check for active data sets. The unmovable table will also be checked. Simulation may indicate that these data sets will move when, in fact, they will not move during a live COMPAKTion.

If you are going to run FASTCPK on a volume with active data sets immediately after a SIMULATION, be sure to specify DSNENQ=TEST or USE on the SIMULATE to get a more accurate prediction of actual results. However, remember that the list of active data sets may change when the operation is run for real, resulting in somewhat different results than the simulation reported.

40.22 COMPAKTOR REPORTS

EXTENTS MAP COMPAKTOR produces a map showing all of the data sets and free spaces on the volume in sequence by location. For SIMULATE and COMPAKT, maps are produced showing the condition of the volume both BEFORE and AFTER the operation. The map shows the data set name, the location and size of each extent, and the extent number and total number of extents. For the first extent of a data set, the map shows the data set organization, the type and quantity of secondary allocation, the last block pointer, the total size of the data set, and how much of the allocated space is unused.

COMPAKTOR normally will highlight the lines for the VTOC and the free spaces, to make them stand out more. (If you plan to look at the map from a terminal, we recommend that you specify HIGHLIGHT=NO, because the highlighting would appear as duplicated lines.)

For non-VSAM data sets, the LAST BLK TTR points to the highest used record in the data set; the first 4 digits are the highest used track (in hex, relative to zero) and the last 2 digits are the highest record number on that track (in hex). EMPTY TRKS is based on the total allocation minus the used tracks for PS and PO data sets.

For ICF VSAM components the LAST BLK TTR contains generated information. The first 4 digits are the last track of the highest used Control Area (in hex, relative to zero) in the cluster as derived from the HI-USED-RBA value in the VVDS entry for the cluster and the last 2 digits are the number of tracks per CA (Control Area), in hex, except the hi-order X'80' flag will be on to indicate that COMPAKTOR was unable to release space on this component. If the TTR contains non-zero data, then EMPTY TRKS will indicate unused tracks in the cluster. These values may be zero for ICF VSAM components in the AFTER MAP.

See the following page for a sample COMPAKTion Extent Map.

For volumes with more than 64K tracks, such as the 3390-9, the map is in a slightly different format with starting and ending cylinders on the left, and 5-digit cylinder numbers. The MAPFORMAT=NEW operand can force this format even on smaller disks.

CONTINUED . . .

40.22 CONTINUED . . .

SAMPLE COMPAKTOR CONTROL CARDS

Using VSR=ALL to release overallocated space in VSAM files and using SIZEKEEP=0 to maximize free space consolidation.

```

● CPK3011 INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P      MESSAGES      DATE 94.274 TIME 08.34.02 PAGE
●
● CPK3051 CARD -- * SIM TYPE=CPK,VOL=CLS005,ACTMESS=NO,VSR=ALL,      *
● CPK3051 CARD -- * SIZEKEEP=0,                                      *
● CPK3051 CARD -- * LOG=YES,CONFMESS=NO                             *

```

BEFORE COMPAKTION

```

● CPK3011 INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P      EXTENTS MAP OF      DATE 94.274 TIME 08.34.02 PAGE 1
●                                     VOLUME CLS005 BEFORE COMPAKTION
●
● START END EXTENT
● TRACK TRACK LENGTH D A T A   S E T   N A M E      D/S SPACE 2ND. LAST BLK TOTAL EMPTY START END
●                                     EXTENT ORG ALLOC ALLOC TTR(HEX) TRKS. TRKS. CC-HH CC-HH
●
● 00000 00000      1 *** IPL AND LABEL RECORDS ***                                0000-00 0000-00
● 00001 00089      89 *** VTOC ***                                01/01                                0000-01 0005-14
● 00090 00134      45 SYS1.VTOCIX.CLS005                                01/01 PS   TRK    0 002C12    45    0 0008-00 0008-14
● 00135 00179      45 SYS1.VVDS.VCLS005                                01/01 EF   TRK    0 002C81    45    0 0009-00 0011-14
● 00180 02204     2025 FLX2378.P.OL.ENROLLMN.D                      01/03 EF   CYL   20 07E80F  5655  3630 0012-00 0146-14
● 02205 05534     3330 FLX2378.P.OL.ENROLLMN.D                      02/03                                0147-00 0368-14
● 05535 05834      300 FLX2378.P.OL.ENROLLMN.D                      03/03                                0369-00 0368-14
● 05835 06254      420 PAS1612.P.TSC.ACTIVITY.D                    01/31 EF   CYL    5 11C08F  4575    30 0389-00 0416-14
● 06255 06764      510 PAS1612.P.TSC.ACTIVITY.D                    02/31                                0417-00 0450-14

```

AFTER COMPAKTION

```

● CPK3011 INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P      EXTENTS MAP OF      DATE 94.274 TIME 08.34.02 PAGE 1
●                                     VOLUME CLS005 AFTER COMPAKTION
●
● START END EXTENT
● TRACK TRACK LENGTH D A T A   S E T   N A M E      D/S SPACE 2ND. LAST BLK TOTAL EMPTY START END
●                                     EXTENT ORG ALLOC ALLOC TTR(HEX) TRKS. TRKS. CC-HH CC-HH
●
● 00000 00000      1 *** IPL AND LABEL RECORDS ***                                0000-00 0000-00
● 00001 00089      89 *** VTOC ***                                01/01                                0000-01 0005-14
● 00090 00134      45 SYS1.VTOCIX.CLS005                                01/01 PS   TRK    0 002C12    45    0 0008-00 0008-14
● 00135 00179      45 SYS1.VVDS.VCLS005                                01/01 EF   TRK    0 002C81    45    0 0009-00 0011-14
● 00180 04754     4575 PAS1612.P.TSC.ACTIVITY.D                      01/01 EF   CYL    5 11C08F  4575    30 0012-00 0316-14
● 04755 07544     2790 FLX4061.P.OL.TRANSACTION.D                    01/01 EF   CYL   20 0AE50F  2790    0 0317-00 0502-14
● 07545 10229     2685 FLX0860.P.OL.ENROLLMN.D                      01/01 EF   CYL   20 0A7C0F  2685    0 0503-00 0681-14
● 10230 12254     2025 FLX2378.P.OL.ENROLLMN.D                      01/01 EF   CYL   20 07E80F  2025    0 0682-00 0816-14
● 12255 14234     1980 FLX1488.P.OL.BALANCE.D                      01/01 EF   CYL   20 07B80F  1980    0 0817-00 0948-14

```

Before COMPAKTion data set PAS1612.P.TSC.ACTIVITY.D was in 31 extents and after COMPAKTion this VSAM data component went into one extent.

40.22 CONTINUED . . .

SUMMARY COMPAKTOR produces a summary showing many statistics for each volume. For SIMULATE and COMPAKT, two columns show the statistics BEFORE and AFTER the operation. There are statistics for the volume as a whole, for various categories of data sets, and for the VTOC. Some uses of the summary map are discussed in [Section 40.15](#).

● CPK3011 INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P		SUMMARY MAP OF VOLUME CLS005	
		BEFORE	AFTER
● VOLUME SUMMARY ---	DEVICE TYPE	3380-K	3380-K
●	NO. OF TRACKS	39825	39825
●	NO. OF CYLINDERS	2655	2655
●	NO. OF TRACKS PER CYLINDER	15	15
●	NO. OF ALLOCATED TRACKS	37889	27184
●	NO. OF ALLOCATED EXTENTS	193	101
●	NO. OF UNMOVEABLE EXTENTS	2	2
●	NO. OF FREE TRACKS	1936	12641
●	NO. OF FREE CYLINDERS	127	842
●	① NO. OF FREE SPACE AREAS	6	1
●	② SIZE OF LARGEST FREE AREA, IN TRACKS	1035	12641
●	PERCENTAGE OF VOLUME IN USE	95	68
●	IBM FRAGMENTATION INDEX	0.257	0.000
● DATASET SUMMARY --	NO. OF EMPTY TRACKS IN PS DATA SETS	0	0
●	NO. OF EMPTY TRACKS IN PO DATA SETS	498	498
●	③ NO. OF EMPTY TRACKS IN VSAM DATA SETS	11960	1255
●	NO. OF PS DATA SETS WITH EMPTY TRACKS	0	0
●	NO. OF PO DATA SETS WITH EMPTY TRACKS	5	5
●	NO. OF UNMOVEABLE DATA SETS	2	2
●	NO. OF DATA SETS WITH 2 OR MORE EXTENTS	11	0
●	NO. OF DATA SETS WITH SUL TRACKS	0	0
●	NO. OF TEMPORARY DATA SETS	0	0
●	NO. OF DATA SETS -- IS ORGANIZATION	0	0
●	NO. OF DATA SETS -- PS ORGANIZATION	1	1
●	NO. OF DATA SETS -- PO ORGANIZATION	5	5
●	NO. OF DATA SETS -- DA ORGANIZATION	0	0
●	NO. OF DATA SETS -- AM ORGANIZATION	0	0
●	NO. OF DATA SETS -- EF ORGANIZATION	94	94
●	NO. OF DATA SETS -- ** ORGANIZATION	0	0
●	NO. OF DATA SETS -- TOTAL	100	100
● VTOC SUMMARY -----	VTOC IS AT RELATIVE TRACK	1	1
●	VTOC SIZE IN TRACKS	89	89
●	NO. OF UNUSED VTOC TRACKS	86	87
●	NO. OF DSCBS PER TRACK	53	53
●	NO. OF MODEL DSCBS	1	1
●	NO. OF FORMAT 0 DSCBS	4605	4614
●	NO. OF FORMAT 1 DSCBS	101	101
●	NO. OF FORMAT 2 DSCBS	0	0
●	NO. OF FORMAT 3 DSCBS	9	0
●	NO. OF FORMAT 4 DSCBS	1	1
●	NO. OF FORMAT 5 DSCBS	1	1
●	NO. OF FORMAT 7 DSCBS	0	0
●	NO. OF DSCBS -- TOTAL	4717	4717
●	NO. OF DSCBS IN USE	112	103
●	PERCENTAGE OF DSCBS IN USE	2	2
●	*** END OF SUMMARY MAP ***		

- ① Reduced 6 free space areas to one.
- ② Largest free area significantly increased.
- ③ Releasing space in over allocated VSAM data sets.

40.22 CONTINUED . . .

COMBINED SUMMARY

COMPAKTOR produces a combined summary, with 2 lines (before and after) with selected statistics from the summary of each volume processed, sorted by volser. Only 1 line is printed for MAP. The combined summary is printed on DDname SYSSUMM, if present, or at the end of SYSMAP if not.

Here is a sample combined summary.

CPK301I INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P										COMBINED SUMMARY				DATE 94.145 TIME 11.32.16		PAGE	1	
VOLSER	DEVTYPE	NUMBER TRACKS	OF DSNS	>1 EXTS	ALLOCATED TRACKS	EXTS	%AL	FREE TRACKS	AREAS	LARGEST	FRAG INDEX	EMPTY VSAM	TRACKS PS	IN PO	VTOC SIZE	%US	TIME (MIN)	COMP CODE
IDPBK0	3380	13275	31	0	1730	32	13	11545	7	5610	0.161	0	0	0	15	4		
	--AFTER-CPK-->		31	0	1730	32	13	11545	2	5773	0.082	0	0	0	15	4	.6	0
IDPLB0	3380	13275	104	9	8084	132	61	5191	31	2651	0.232	16	2	1367	15	15		
	--AFTER-CPK-->		104	0	8084	105	61	5191	1	5191	0.000	16	2	1367	15	15	1.1	0
IDPLB1	3380-K	39825	365	38	32251	493	81	7574	48	2051	0.243	48	416	5242	15	48		
	--AFTER-CPK-->		365	3	32251	366	81	7574	1	7574	0.000	48	416	5242	15	47	3.9	0
IDPLB5	3380-K	39825	458	49	38650	615	97	1175	82	130	0.615	668	160	4422	15	61		
	--AFTER-CPK-->		458	2	33540	462	84	6285	2	5988	0.043	50	0	90	15	58	5.9	0
IPLXA2	3380	13275	117	0	9557	118	72	3718	2	2002	0.095	2307	0	1397	15	15		NO
	--AFTER-CPK-->		117	0	9557	118	72	3718	2	2002	0.095	2307	0	1397	15	15	.2	IMPRV
SMSLB0	3380	13275	351	3	10279	364	77	2996	32	1650	0.213	2532	374	53	15	45		
	--AFTER-CPK-->		351	0	10279	352	77	2996	2	1503	0.098	2532	374	53	15	45	2.9	0
TSOWK1	3380	13275	249	7	7591	268	57	5684	31	2205	0.178	0	119	268	10	48		
	--AFTER-CPK-->		249	0	7482	141	56	5793	2	2909	0.090	0	119	268	10	48	1.6	0
CPK RELEASED 11382 TRKS AND ELIMINATED 602 FREE SPACE AREAS FROM 37 VOLUMES WITH 685736 TRKS (44% WAS ALLOCATED)																		
*** END OF SUMMARY MAP ***																		

Many of these fields are extracted from the SUMMARY REPORT on the previous page, and are described in [Section 40.15](#). Some are unique to this report. Briefly, the fields are:

NUMBER OF TRACKS	Total tracks on the volume
NUMBER OF DSNS	Total data sets on the volume
>1 EXTS	Number of data sets with more than one extent
ALLOCATED TRACKS	Total allocated tracks
ALLOCATED EXTS	Total allocated extents
ALLOCATED %AL	Percentage of total tracks allocated
FREE TRACKS	Total tracks not allocated
FREE AREAS	Number of free space areas
FREE LARGEST	Largest free area size in tracks
FRAG INDEX	IBM fragmentation index
EMPTY TRACKS IN..	Total unused tracks in VSAM, PS, and PO data sets
VTOC SIZE	Total size of the VTOC in tracks
VTOC %US	Percentage of the VTOC DSCBs in use
TIME	CPK actual elapsed time in minutes (and tenths if less than 99 minutes). For simulation of Fast COMPAKTion (SIM TYPE=FASTCPK) only, this will be a conservative estimate; the actual elapsed time will usually be less.

40.22 CONTINUED . . .

COMP CODE

CPK completion on this volume:

0 - successful COMPAKTion

ERROR - an error occurred, look for error messages

NO IMPRV - unable to improve the volume

BYPASS - conditional keywords bypassed volume

IN RECVY - FASTCPK previously failed; rerun CPK

40.23 COMPAKTOR EXAMPLES

The following examples illustrate the most common ways of executing COMPAKTOR. Note that for convenience, all STEPLIB/JOBLIB DD statements have been omitted in the examples; they may be required, depending on your installation's placement of COMPAKTOR.

EXAMPLE 1 MAP a direct access volume using a DISK DD statement.

```
//MAP          EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD    SYSOUT=*
//SYSMAP       DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//DISK1        DD    UNIT=3380,VOL=SER=MYDISK,DISP=OLD
//SYSIN        DD    *
                MAP
```

EXAMPLE 2 MAP all online disk volumes.

```
//MAP          EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD    SYSOUT=*
//SYSMAP       DD    SYSOUT=*
//SYSSUMM     DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//SYSIN        DD    *
                MAP VOL=*
```

EXAMPLE 3 MAP a disk volume dumped by FDR, from the backup tape.

NOTE: COMPAKTOR can only map a full-volume backup tape. To map a data set or full backup, including VSAM information, use FDRABRP "PRINT TVTOC" .

```
//MAP          EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD    SYSOUT=*
//SYSMAP       DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//TAPE1        DD    UNIT=TAPE,VOL=SER=FDRDMP,DISP=OLD,
//              DSN=FDR.DUMP
//SYSIN        DD    *
                MAP      FROMDD=TAPE1
```

EXAMPLE 4 SIMULATE doing a COMPAKT-from-backup on selected online disk volumes. All online volumes whose serial begins with TSO or PROD will be simulated, using SIZEKEEP=0.

```
//SIM          EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD    SYSOUT=*
//SYSMAP       DD    SYSOUT=*
//SYSUDUMP     DD    SYSOUT=*
//SYSIN        DD    *
                SIMULATE TYPE=FASTCPK,VOL=(TSO*,PROD*),SIZEKEEP=0
```

NOTE: Innovation recommends using FASTCPK instead, but if you are planning on COMPAKTING volumes with a backup tape, then you should specify SIZEKEEP=0 to achieve maximum free space consolidation and data set extent reduction.

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 5 SIMULATE doing a Fast COMPAKTion on selected online disk volumes. All online volumes whose serial begins with TSO or PROD will be simulated, using all Fast COMPAKTOR defaults. Only summaries are printed.

```
//SIM          EXEC   PGM=FDRCPK,REGION=4M
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM      DD     SYSOUT=*
//SYSIN        DD     *
SIMULATE      TYPE=FASTCPK,VOL=(TSO*,PROD*),MAPS=SUMMARY
```

EXAMPLE 6 SIMULATE Fast COMPAKTion on a volume and attempt to reduce the number of free areas to as few as possible.

```
//SIM          EXEC   PGM=FDRCPK,REGION=4M
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM      DD     SYSOUT=*
//SYSIN        DD     *
SIM TYPE=FASTCPK,VOL=TS0001,SIZEKEEP=(100,100,60)
```

EXAMPLE 7 SIMULATE Fast COMPAKTion with space release and the default of SIZEKEEP=(100,90,60) is recommended for volumes with many multi-extent data sets or with an extremely large number of small data sets (especially 3390-9). This will simulate the result of releasing space and COMPAKting.

```
//SIM          EXEC   PGM=FDRCPK,REGION=4M
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
SIM TYPE=FASTCPK,VOL=DB*,DSNENQ=USE,PSRLSE=ALL,
PORLSE=ALL,VSRLSE=ALL,CPKFRAG1=2,CPKFREEX=20,
NOSECOND=NORLSE
```

HINTS:

1. To run these simulations for real, just change SIM to CPK.
2. Add LOG=YES to real COMPAKTions to display FDRW82 operator messages when COMPAKTOR begins and ends modifying each volume.
3. If you plan to run a real COMPAKTion immediately after a SIMULATE, be sure to add DSNENQ=USE to the simulation to get a more ACCURATE REPORT of what FASTCPK will do.

EXAMPLE OF COMBINED SUMMARY REPORT FROM SIMULATION

VOLSER	DEVTYPE	TRACKS	DSNS	EXTS	TRACKS	EXTS	%AL	TRACKS	AREAS	LARGEST	INDEX	VSAM	PS	PO	SIZE	%US	(MIN)	COMP
IDPLB1	3380-K	39825	365	38	32251	493	81	7574	48	2051	0.243	48	416	5242	15	48	(EST)	
-- AFTER-CPK -->			365	3	32251	366	81	7574	1	7574	0.000	48	416	5242	15	47	3.9	0

Multi-Extents Data Sets

Before 38

After 3

Free Areas

Before 48

After 1

Estimated Elapsed Time

3.9 minutes

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 8 Space Release and Fast COMPAKTion on multiple volumes. The first step will release free space from all inactive PS, PO and VSAM data sets. The second step will do a FASTCPK on any of those volumes where there are more than 20 free space areas. Active data sets will not be moved. This 2-step technique is slightly faster than doing a single-step FASTCPK including space release, since the FASTCPK step will move fewer tracks than the single-step job.

```
//CPKRLSE      EXEC   PGM=FDRCPK,REGION=4M
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM      DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
      COMPACT TYPE=RLSE,VOL=(TSO*,TEST*),
              PSRLSE=ALL,PORLSE=ALL,VSRLSE=ALL,NOSECOND=NORLSE
//FASTCPK      EXEC   PGM=FDRCPK,REGION=4M,COND=EVEN
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM      DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
      COMPACT TYPE=FASTCPK,VOL=(TSO*,TEST*),
              CPKFREEEX=20,LOG=YES,UNABLE=IGNORE
```

EXAMPLE 9 Fast COMPAKTion with selective space release. The specified ICF VSAM data sets will have space released (leaving 20% free space) if they have secondary allocation specified. Data sets CICS.TEST.** will not be moved.

```
//FASTCPK      EXEC   PGM=FDRCPK,REGION=4M
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM      DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
      COMPACT TYPE=FASTCPK,STORGRP=TEST*,NOSECOND=NORLSE,CPKFREEEX=10
      SELECT DSN=TEST.**,DSORG=EF,VSRLSE=ALL,%FREE=20
      SELECT DSN=CICS.TEST.**,POS=KEEP
```

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 10 COMPAKT a SMS-managed disk volume from a FDR dump tape. The dump was performed at an earlier time. Since the disk dumped to the backup tape was SMS-managed, COMPAKTOR forces CPYVOLID=YES. The output volume specified by the VOL= parameter must be SMS-managed.

This example will also work if both the backup and the target disk are NOT SMS-managed. It can also be used when the output disk is a different capacity than the original (e.g., 3390-3 to 3390-2) as long as all the data on the original disk will fit on the output.

```
//COMPAKT      EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT      DD    SYSOUT=*
//SYSMAP        DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//TAPE1         DD    DSN=FDR.BACKUP.VSMS001,DISP=OLD
//SYSIN         DD    *
                COMPAKT  VOL=SMS001,SIZEKEEP=0
```

NOTE: SIZEKEEP=0 should always be used whenever a backup tape is used so that COMPAKTOR will maximize free space and merge all multi-extent data sets.

EXAMPLE 11 COMPAKT-from-backup an active SMS-managed disk volume using dynamic allocation of the disk. First the disk is dumped and then restored in COMPAKTed form. The data sets will be enqueued for the duration of the DUMP and COMPAKTion. Two tape drives are used to minimize tape mount and rewind delays. Console message FDRW82 will be issued at the beginning and end of the COMPAKTOR restore. SIZEKEEP=0 is specified to force COMPAKTOR to make all data sets movable, so that all free space will be collected into 1 or 2 areas and all multi-volume data sets are (usually) merged into single extents; this may be useful in cases when FASTCPK cannot accomplish this.

This example will also work if disk is NOT SMS-managed.

```
//COMPAKT      EXEC  PGM=FDRCPK,REGION=2048K
//SYSPRINT      DD    SYSOUT=*
//SYSMAP        DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//TAPE1         DD    DSN=FDR.DUMP.VSMS123,DISP=(,KEEP),
//              UNIT=(3480,2),VOL=(, , ,99)
//SYSIN         DD    *
                DUMP    TYPE=FDR,DSNENQ=USE,COMPRESS=ALL
                CPK      VOL=SMS123,DUMP=YES,ACTMESS=NO,LOG=YES,
                        SIZEKEEP=0
```

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 12 COMPAKT a disk volume onto another disk volume. First the disk volume is dumped, and then restored onto the target volume. The volume serial of the target volume will be changed to OLDDISK on completion of the COMPAKTion. At that time, the original volume should be VARYed offline, and the new volume should be MOUNTed. If the original volume had an indexed VTOC, it must be rebuilt on the new volume with ICKDSF.

```
//COMPAKT      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=A
//SYSMAP       DD     SYSOUT=A
//SYSUDUMP     DD     SYSOUT=A
//TAPE1        DD     DSN=FDR.DUMP.VOLDDSK,DISP=(,KEEP),
//              UNIT=3480,VOL=(,,99)
//DISK1        DD     UNIT=3380,VOL=SER=OLDDISK,DISP=OLD
//NEWDISK1     DD     UNIT=3380,VOL=SER=NEWDISK,DISP=OLD
//SYSIN        DD     *
                DUMP   TYPE=FDR,COMPRESS=ALL
                CPK     DUMP=YES,TODD=NEWDISK1
```

EXAMPLE 13 RELEASE space from selected online disk volumes. All online volumes whose serial begins with TSO will be released. COMPAKTOR will attempt to release all excess tracks or cylinders from all PS, PO and ICF VSAM data sets, leaving 10% of the allocated space as free space. However, space will not be released from data sets which do not have a secondary allocation quantity

ICF VSAM clusters with IPCS in their name will not have space released. Active data sets will not be released.

```
//RELEASE      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSSUMM     DD     SYSOUT=*
//SYSIN        DD     *
                COMPAKT TYPE=RLSE,VOL=TSO*,PSRLSE=ALL,PORLSE=ALL,VSRLSE=ALL,
                %FREE=10,NOSECOND=NORLSE
                SELECT  DSN=**IPCS**,VSRLSE=NO
```

EXAMPLE 14 Convert a 3390-2 to a 3390-3 or RAMAC DASD volume with COMPAKTOR, moving the VTOC to the first third of the volume. The VTOC index and the VVDS (if they exist) will be placed adjacent to the VTOC (if they do not exist, it will not be considered an error). The volume serial of the output disk will be changed to that of the original volume, and the new volume will automatically be placed offline at the end of the COMPAKTion. At that time, the original volume should be VARYed offline, and the new volume should be MOUNTed. If the original volume had an indexed VTOC, it must be rebuilt on the new volume with ICKDSF.

```
//COMPAKT      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//DISK1        DD     UNIT=3390,DISP=OLD,
//              VOL=SER=D33902                                <--- 3390-2
//NEWDISK1     DD     UNIT=3390,DISP=OLD,
//              VOL=SER=D33903                                <--- 3390-3
//TAPE1        DD     UNIT=3490,DISP=(,CATLG),
//              DSN=BACKUP.D33902,VOL=(,,20)
//SYSIN        DD     *
                DUMP   TYPE=FDR,COMPRESS=ALL
                CPK     DUMP=YES,TODD=NEWDISK1,VTOC=COMPAKT,SIZEKEEP=0,
                OVERRIDE=YES,SELTERR=NO
                SELECT  ***VTOC,POS=11130000,SIZE=45
                SELECT  DSN=SYS1.VTOCIX.***,POS=VTOC
                SELECT  DSN=SYS1.VVDS.***,POS=VTOC
```

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 15 Convert a 3390-3 to a 3390-2 volume with COMPAKTOR, moving the VTOC to the beginning of the volume. The VTOC index and the VVDS (if they exist) will be placed adjacent to the VTOC. The volume serial of the output disk will be changed to that of the original volume, and the new volume will automatically be placed offline at the end of the COMPAKTion. At that time, the original volume should be VARYed offline, and the new volume should be MOUNTed. If the original volume had an indexed VTOC, it must be rebuilt on the new volume with ICKDSF. This procedure will work only if there are no unmovable data sets in the upper third of the 3390-3 and if the total allocated tracks do not exceed the capacity of the output 3390 data set.

```
//COMPAKT      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT      DD     SYSOUT=*
//SYSMAP        DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//NEWDISK1      DD     UNIT=3390,DISP=OLD,
//              DD     VOL=SER=D33902                      <---3390-2 output
//TAPE1         DD     UNIT=(3490,2),DISP=(,CATLG),
//              DD     DSN=BACKUP.D33903,VOL=(,,99)
//SYSIN         DD     *
DUMP            TYPE=FDR,DATA=USED
CPK            DUMP=YES,TODD=NEWDISK1,VTOC=COMPAKT,SIZEKEEP=0,
               OVERRIDE=YES,SELTERR=NO,VOL=D33903          <---3390-3 input
SELECT         ***VTOC,POS=BEGIN
SELECT         DSN=SYS1.VTOCIX.**,POS=VTOC
SELECT         DSN=SYS1.VVDS.**,POS=VTOC
```

EXAMPLE 16 COMPAKT a disk in place from a backup (DUMP=YES). This example can be used instead of FASTCPK if you have a need to position a large number of data sets (positioning increases the elapsed time of FASTCPK). All DB2 clusters are to be moved next to the VTOC (note that the second level of DSNDBD designates the data component name for DB2) and all other VSAM will have enough unused CAs released so that 25% of the remaining CAs are free for expansion.

```
//COMPAKT      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT      DD     SYSOUT=*
//SYSMAP        DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     UNIT=(3490,2),DISP=(,CATLG),
//              DD     DSN=BACKUP.D33901,VOL=(,,99)
//SYSIN         DD     *
DUMP            TYPE=FDR,DATA=USED,DSNENQ=USE
CPK            DUMP=YES,VOL=D33901,SIZEKEEP=0
SELECT         DSN=*.DSNDBD.*.*.10001.A+++,POS=VTOC
SELECT         ALLDSN,DSORG=EF,VSRLSE=ALL,%FREE=25
```

CONTINUED . . .

40.23 CONTINUED . . .

EXAMPLE 17 Expand the VTOC on a volume. This can only be done if you COMPAKT a disk from a backup; this example uses DUMP=YES to do so. If necessary, COMPAKTOR will move data sets in order to free space to expand the VTOC. The SIZE= operand on the SELECT statement specifies the new VTOC size in tracks. If you added the POS= operand you could also move the VTOC. The other SELECT statements will move the VTOCIX and VVDS next to the VTOC.

```
//COMPAKT      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//TAPE1        DD     UNIT=CART,DISP=(,CATLG),
//              DSN=BACKUP.D33901,VOL=(,,99)
//SYSIN        DD     *
DUMP           TYPE=FDR,DATA=USED,DSNENQ=USE
CPK            DUMP=YES,VOL=D33901,SIZEKEEP=0,VTOC=COMPAKT,OVERRIDE=YES
SELECT        ***VTOC,SIZE=30
SELECT        DSN=SYS1.VTOCIX.*,POS=VTOC
SELECT        DSN=SYS1.VVDS.*,POS=VTOC
```

EXAMPLE 18 Release unused space within PS and PO data sets on all online volumes starting with "TSO". CPK will not COMPAKT the volumes, only release space. PS data sets with a group name of TEST will be completely free. All other data sets will be left with 10% free space.

```
//RELEASE      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
CPK            TYPE=RLSE,VOL=TSO*,PSRLSE=ALL,PORLSE=ALL,%FREE=10
SELECT        DSN=TEST**,DSORG=PS,RLSE=TRK
```

EXAMPLE 19 Release all free cylinders in data sets beginning with "A.B.C", and reduce all data sets (including ICF VSAM clusters) beginning with "C.D" to contain 20% free tracks, on all online volumes. Note that the RLSE= and VSRLSE= operands appear only on SELECT statements, and not on the CPK statement; this is the proper way to release space ONLY on selected data sets.

```
//RELEASE      EXEC   PGM=FDRCPK,REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSMAP       DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SYSIN        DD     *
CPK            TYPE=RLSE,VOL=*
SELECT        DSN=A.B.C,RLSE=ALL
SELECT        DSN=C.D.*,%FREE=20,RLSE=TRK,VSRLSE=ALL
```

CONTINUED . . .

40.24

COMPAKTOR ISPF DIALOGS

If you have installed COMPAKTOR using the ISPF interactive install procedure you have available easy-to-use ISPF panels to:

- a. MAP volumes
- b. SIMULATE FASTCPK and TYPE=RLSE jobs
- c. Update the CPK unmovable table
- d. Change COMPAKTOR defaults in the Global Option Table

Select option "C" COMPAKTOR in the FDR Primary Options Menu, as illustrated on the following figure:

PANEL A.C

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- FDR PRIMARY OPTIONS MENU -----
OPTION ==> C_

      1  REPORTS      -  ABR REPORTING FUNCTIONS
      2  RESTORE      -  ABR DATA SET RESTORE
      3  ARCHIVE      -  ABR DATA SET ARCHIVE OR SUPERSCRATCH
      4  BACKUP       -  ABR DATA SET BACKUP
      5  REMOTE Q     -  ABR REMOTE QUEUE UTILITY FUNCTIONS

      C  COMPAKTOR    -  COMPAKTOR MAP AND SIMULATION REPORTS
      R  RELEASE      -  COMPAKTOR RELEASE
  
```

**SAMPLE
COMPAKTOR
MAP AND
SIMULATION
PANEL**

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- COMPAKTOR MAP/SIMULATION PANEL -----

SELECT OPTION      ==>                                | PROCESSING OPTION ==>
BLANK  -  MAP                               | BLANK - FOREGROUND
F      -  SIMULATE FASTCPK                   | B    - BATCH
S      -  SIMULATE COMPAKTION                |
R      -  SIMULATE RELEASE                   | TEST FOR ACTIVE DSN==> NO (YES, NO)

VOLUME SERIAL(S) ==>
OR STORAGE GROUP(S) ==>

PS RELEASE OPTION ==> %FREE ==>                | MERGE VSAM EXTENTS ==> YES (YES, NO)
PO RELEASE OPTION ==> %FREE ==>                |
VSAM RELEASE OPTION ==> %FREE ==>              | SIZEKEEP ==>( 100 , 90 , 60 )
BLANK  -  DO NOT RELEASE                     |
A      -  ALL UNUSED                         | CPKFREEX ==> 0
R      -  ROUND (PS/PO)                     |
T      -  ALL UNUSED TO TRACK (PS/PO)        | REPORT FORMAT      ==>          (C=CHANGE)
NOSECOND ==> NORLSE (RLSE, NORLSE)           | PRINT BOTH BEFORE AND AFTER MAPS
                                           | PRINT EXTENT MAP AND SUMMARY REPORT
                                           | REPORT TO SCROLL DATA SET
                                           | HIGHLIGHTING IS NOT REQUIRED
                                           | CREATE 58 LINES PER PAGE
  
```

**COMPAKTOR
INSTALLATION
OPTIONS**

See [Section 90.14](#) for the ISPF panels used to set COMPAKTOR global options and defaults. See [Section 90.30](#) for the ISPF panels used to maintain the COMPAKTOR Unmovable Table.

CONTINUED ...

40.25 CPK HINTS AND RECOMMENDATIONS



- **I ran FASTCPK on a 3390-3 and it reduced my free areas from 322 to 26.47 data sets were in more than one extent before FASTCPK, and 19 after FASTCPK. FASTCPK ran in 2.9 minutes, which is great, but I want even better free space consolidation and to eliminate all my multi-extent data sets.**

FASTCPK's default is to reduce the number of free space areas by 90%.

If you want COMPAKTOR to do better consolidation, you may specify the SIZEKEEP option. The default for SIZEKEEP is (100,90,60). For maximum consolidation, specify SIZEKEEP=(0,100,1); then FASTCPK will attempt to consolidate all of the free space into 2 or 3 free space areas and will consolidate all data sets with more than one extent into one extent. However, the additional consolidation will require some additional run time.

HINT: Innovation recommends that you run with the default to minimize elapsed time, and specify SIZEKEEP=(0,100,1) for better free space and extent consolidation or SIZEKEEP=(100,100,60) for better space consolidation. For 3390-9 volumes, you should always use SIZEKEEP=(100,90,60), the default.

► **Recommended Options**

Average FASTCPK time is under 4 minutes for 3390-3. On most volumes the average time will actually be 3 minutes.

For the best performance we actually recommend that you run free space release

(CPK TYPE=RLSE) as a separate step just prior to the FASTCPK run. For example,

First Step:

```
COMPAKT TYPE=RLSE , VOL=TSO* , NOSECOND=NORLSE ,
PSRLSE=TRK , PORLSE=ALL , VSRLSE=ALL , %POFREE=10 , %VSFREE=10
```

Second Step:

```
COMPAKT TYPE=FASTCPK , VOL=TSO* , LOG=YES , CPKFRAGI=3 ,
CPKFREEEX=10 , UNABLE=IGNORE , SIZEKEEP=( 100 , 90 , 60 )
```

- Activity level varies daily on storage volumes. It is hard to predict. As a general rule for FASTCPK during the week you might want to COMPAKT only those volumes with a CPKFRAGI of 3 (fragmentation index of .300) or more or those with at least 10 free space areas.
- Specifying LOG=YES will cause COMPAKTOR to issue FDRW82 COMPAKTOR REORGANIZATION STARTED VOL=vvvvvv and FDRW82 COMPAKTOR REORGANIZATION SUCCESSFUL VOL=vvvvvv messages on the console. These messages will aid the operator in determining what volumes are being processed.
- Significant amounts of space are wasted within PS, PO and VSAM data sets including DB2. FASTCPK can release all or a percentage of the space during FASTCPK. To preserve some free space within data sets, specify %FREE= or %PSFREE=, %POFREE=, %VSFREE=. It is recommended that some free space be left in VSAM and PO files.
- NOSECOND=NORLSE should be specified so data sets without secondary allocations will not be eligible for space release.
- UNABLE=IGNORE will suppress a U0888 abend if COMPAKTOR encounters an unusual circumstance like all of the

CONTINUED . . .

40.25 CONTINUED . . .

► **How do I avoid a U0888 abend if COMPAKTOR gets a CPK554E UNABLE TO COMPAKT THIS VOLUME REASON=ALGORITHMS FAILED-VOLUME NOT CHANGED**

In some unusual circumstances, FASTCPK might not be able to COMPAKT a volume (all of the data set placement algorithms have failed). This type of error will result in the general ABEND code U0888 at the end of the job. Since the volume was not changed, you can code UNABLE=IGNORE on the COMPAKT statement to suppress the U0888 abend for this type of error.

► **How accurate is the elapsed time estimate in the SIMULATION COMBINED SUMMARY report?**

Very accurate. The COMPAKTOR survey results and our own testing showed that over 95% of the FASTCPK jobs ran at or less than the estimated time. Our experience on EMC, RAMAC and RVA have shown even better results, since these new devices have even faster access time and make more efficient use of cache.

HINT: If you are going to run FASTCPK on a volume with active data sets, right after a SIMULATION, add DSNENQ=USE to the SIM TYPE=FASTCPK run to get a more accurate report on what FASTCPK will do.

► **Currently we run COMPAKTOR under ABR by specifying DUMP TYPE=FDR, COMPAKT on the weekly full volume backups. What is the best way for us to take advantage of FASTCPK?**

INNOVATION recommends that you run FASTCPK in separate job(s) before the ABR full volume backups. FASTCPK, with its built-in recovery capability, does not require a prior backup. FASTCPK will run faster than COMPAKting from the backup tape under ABR.

► **I received a message “CPK554E UNABLE TO COMPAKT THIS VOLUME REASON=RECOVERY LOG FULL”. What does this mean?**

If a significant number of tracks need to be moved, it is possible that the CPK recovery log will fill up and COMPAKTOR will go into recovery and backout the changes. If you have encountered this situation contact INNOVATION for assistance.

► **How many volumes can be processed by FASTCPK at the same time?**

It varies significantly between users. FASTCPK is I/O intensive so it depends on the number and type of paths to your DASD, and on the other types of processing that will be occurring concurrently. The average is about 10 volumes at a time with some large shops doing up to 30 at a time. One user runs 31 volumes at the same time and they all complete within 11 minutes. You should experiment to see what works best in your configuration.

A given FASTCPK job processes only one volume at a time, so to run concurrently you must set up multiple jobs. For best results, put volumes on the same DASD string into the same job.

40.25 CONTINUED . . .

► **My VVDS is in multiple extents.
Can COMPAKTOR merge a VVDS into one extent?**

As a default, COMPAKTOR considers a VVDS as unmovable. Many times, a VVDS will go into extents and this can impact overall DASD performance.

If you specify OVERRIDE=YES and SELECT DSN=SYS1.VVDS.Vxxxxxx,POS=VTOC, COMPAKTOR will merge all the extents of the VVDS into one extent and position the VVDS next to the VTOC.

This type of operation can be done only when the VVDS is not in use and requires COMPAKTOR to take a backup (TYPE=CPK,DUMP=YES). The safest procedure is:

1. VARY the volume OFFLINE to all other systems with shared DASD before COMPAKTOR starts.
2. Do not specify ACTMESS=NO, and do not reply IGNORE if message FDRW81 appears informing you that there are open data sets on the volume.
3. VARY the volume back ONLINE to sharing systems after COMPAKTOR completes.

► **We have the need to run FASTCPK while the volumes are in use. What impact will this have on my users?**

FASTCPK by default, will automatically detect active data sets. Linklist data sets on the system FASTCPK is run on are automatically made unmovable by FASTCPK. Other active data sets not ENQ'ed by the operating system should be put in COMPAKTOR unmovable table (a list of such data sets appears in [Section 40.21](#) of the manual). FASTCPK will issue an ENQ/RESERVE on the VTOC for the duration of the COMPAKTion. This will prevent users on the system where COMPAKTOR is running from allocating, scratching or obtaining secondary extents. Data sets that are in use at the time FASTCPK is started will still be available to the end user. On other systems with shared DASD, the entire volume will be unavailable while it is being COMPAKted. Shortly after FASTCPK starts, FASTCPK will DEQ data sets which are not going to be moved, and those data sets will be available again to end users on the system where COMPAKTOR is running. If any jobs receive message IEF211I while CPK is running, contact Innovation for assistance.

HINT: FASTCPK will not ENQ data sets in the unmovable table. So if you want certain data sets always accessible on a volume that you are COMPAKting, put those data sets in the COMPAKTOR unmovable table. For VSAM data sets, the component name is specified in the COMPAKTOR unmovable table.

► **TSO users receive the following messages during logon for the ISPF profile data set if FASTCPK is COMPAKting that volume and has the ISPF profile data set ENQed:**

```
IKJ56225I DATA SET dsn ALREADY IN USE, TRY LATER+
IKJ56225I DATA SET IS ALLOCATED TO ANOTHER JOB OR USER
```

Without the ISPPROF data set, they cannot access ISPF. This can be circumvented by adding a generic filter mask to the CPK unmovable table for the ISPF profile data sets. Use FDRZAPOP or the ISPF Panels A.I.5 to update the CPK unmovable table with (CPKUNMOV DSN=*.ISPPROF). COMPAKTOR will not ENQ on data sets that are in the CPK unmovable table.

40.25 CONTINUED . . .

► **Occasionally we see free space errors in the VTOC. Can COMPAKTOR correct free space errors?**

COMPAKTOR automatically corrects VTOC errors during any COMPAKTion or RELEASE run.

A simple RELEASE run of one data set runs in less than a minute and will automatically rebuild the VTOC, including the free space descriptors.

Example:

```
CPK TYPE=RLSE,VOL=xxxxxx  
SELECT DSN=MY.DATA SET,RLSE=ALL
```

► **What happens if the operator cancels COMPAKTOR before it is completed?**

Since the majority of interrupted COMPAKTOR jobs result from operators issuing CANCEL commands, COMPAKTOR now includes a CANCEL protection function. If a CANCEL is issued against a COMPAKTOR job at any time during the restore phase (while the VTOC, VVDS, or data tracks are actually being updated or moved), COMPAKTOR will intercept that CANCEL and issue an FDRW99 message to the operator indicating that the CANCEL will result in volume corruption, and giving the operator the option of ignoring the CANCEL, cancelling the job, or bringing COMPAKTOR to a graceful termination when it finishes processing the current volume.

The operator should always reply S (Stop) so that COMPAKTOR will just complete the volume currently in progress before terminating.

HINT: You may want to automate the response to the FDRW99 messages to always be S.

► **What happens if the system crashes during a FASTCPK?**

You can resubmit the COMPAKTOR job that was running and FASTCPK will automatically recover the volume. If you do not know which volume was being processed submit a MAP job with the command MAP VOL=*,MAPS=SUMMARY.

COMPAKTOR will print a CPK586W message for each volume on which you must resubmit FASTCPK to recover the volume.

► **How long does it take for FASTCPK to recover a volume?**

FASTCPK recovery will take about 2 times as much elapsed time as COMPAKTOR spent COMPAKTing the volume before the failure (3390-9 will take longer). For example, if the FASTCPK was processing the volume for 3 minutes, recovery will be completed in 6 more minutes.

We have had some interesting calls on FASTCPK recovery. Some users noticed happily that if FASTCPK encounters an I/O error writing a track of data that it will automatically go into recovery and return the volume to its original status.

40.25 CONTINUED . . .

Simple Conversion

...Minutes to Convert JCL and Control Cards

IBM DEFRAG			Fast COMPAKTOR		
//STEP 4	EXEC	PGM=ADRDSSU	SIMPLE CONVERSION	//FASTCPK	EXEC PGM=FDRCPK
//SYSPRINT	DD	SYSOUT=*		//SYSPRINT	DD SYSOUT=*
//DASD1	DD	VOL=		//SYSMAP	DD SYSOUT=*
//	DD	UNIT=		//SYSIN	DD *
//SYSIN	DD	*		COMPAKT TYPE=FASTCPK,VOL=TS0001	
DEFRAG	DDNAME(DASD1)				

SAMPLE FASTCPK USER EXPERIENCES

# of VOL	Type	AVG # of DSN	Free Area		% Reduction In Free Area	Multi-Extent DSN		AVG Elapsed Time (M.S.)
			Before	After		Before	After	
10	3380-E	1061	4067	20	99%	363	20	5.6
8	3380-K	1145	3387	16	99%	298	5	5.4
3	3380-E	1051	1165	6	99%	96	0	4.7
2	3390-2	1127	892	3	99%	78	0	4.4
10	3380-E	208	860	19	97%	91	0	5.7
4	3390-2	258	95	2	97%	14	0	4.7
10	3390-3	71	336	35	89%	44	32	3.2
4	3390-3	67	114	11	90%	18	12	2.7
10	3390-3	71	353	39	89%	35	24	1.9
10	3390-3	76	340	33	90%	47	37	3.6

Totals			Total # of Free Area's		Total % Reduction	Total # of Multi-Extent DSN		Avg Elapsed Time
			Before	After	In Free Area	Before	After	
391 volumes			32254	1314	96%	3357	785	4.5 minutes
			A 76% Reduction					

The results above are from a typical user. They COMPAKTed 391 DASD volumes. The volumes were separated into three pools:

1. A large data set pool with fewer than 100 data sets per volume.
2. A medium pool with about 300 data sets per volume.
3. A small pool with over 1000 data sets per volume.

The average elapsed time was 4.5 minutes and the number of free areas were reduced by 96%.

40.26 BENEFITS OF FASTCPK ON VIRTUAL DISK ARRAYS

Some customers have asked what use there will be for COMPAKTOR on virtual disk arrays such as the IBM RVA (RAMAC Virtual Array), the StorageTek Iceberg and the StorageTek SVS (Shared Virtual Array). These subsystems write output tracks to new contiguous locations, and have internal procedures and algorithms to consolidate tracks that were written previously. Therefore, internally, this device does not waste any space for over-allocated data sets, and consolidating extents of multi-extent data sets will not provide a performance benefit.

Externally, though, as seen by the Operating System, the virtual disk array still appears to be a set of conventional DASD volumes. Each logical volume has a fixed number of cylinders, and space is managed in the usual way with a VTOC and a VTOC index. If a data set is over-allocated, then the unused space is not available for allocation to other data sets. If the free space on a (logical) volume is fragmented, then it becomes difficult to allocate new data sets; the allocation may fail, or the data set may be allocated in multiple extents. Secondary allocations also may fail or be broken up onto multiple extents. The total number of extents for a data set is still limited to 16 per volume (except for VSAM, PDSE, HFS, and PSE), so if a primary or secondary allocation is obtained in multiple extents, the data set will not be able to get the full amount of space that the user intended when he coded the SPACE parameter. Jobs may get SB37 ABENDs. If you have STOPX37 or a comparable product that avoids Sx37 ABENDs by allowing the data set to be extended to another volume, then you may end up with many multi-volume data sets.

COMPAKTOR solves these problems by consolidating the free space on a volume, by merging multi-extent data sets into a single extent, and by releasing unused space.

If a volume is SMS-managed, then SMS will stop allocating new data sets on the volume when it reaches its high threshold. For example, if the (logical) volume is a 3390-3 with 50,085 tracks, and the high threshold for the storage group is 80%, and a user allocates a data set with 40,000 tracks, then SMS will consider the volume to be full, and will send new allocations to other volumes in the storage group. However, if only 30,000 tracks are used in the 40,000 track data set, you can use FASTCPK to release the over-allocated space. The volume will then be only 60% allocated, and SMS will be able to use it again for new data sets.

Average run time for FASTCPK on a virtual disk array to consolidate free space, merge extents, and release unused space is about 2 minutes. If you just want to release unused space, TYPE=RLSE usually runs in less than 20 seconds.

However, COMPAKTOR can also use the optional Snapshot feature available on the virtual disk arrays to reduce the elapsed time of a FASTCPK to seconds. To use it, your hardware must have the Snapshot feature installed and you must also be licensed for FDR InstantBackup ([see Section 26](#)). To invoke it, code the SNAPSHOT=YES operand on the COMPAKT TYPE=FASTCPK statement.

50.01 FDRABR VOLUME BACKUPS

ABR Program FDRABR (Automatic Backup Restore) automates the execution of FDR full-volume and data set backups and restores for the purposes of:

- **data availability** - creating backup copies of DASD data sets to protect against physical loss or logical damage.
- **space management** - identifying data sets which do not need to be on DASD (usually based on the last date they were used) and moving them to backups (called "archiving"), from which they can be automatically recalled if needed. Data sets which will never be needed again can be scratched without creating a backup.
- **disaster recovery** - creating backups from which all or part of your DASD data can be quickly recreated at a disaster recovery site.
- **tape mount management (TMM)** - maximizing tape volume usage by staging small tape data sets to a disk buffer, then moving many of them to a tape volume.

These objectives are accomplished by the use of several ABR functions:

- This section describes ABR Volume Backups (DUMP TYPE=FDR/ABR/AUTO/DSF) which are used for "data availability" at the volume level.
- ABR ARCHIVE (DUMP TYPE=ARC) and SUPERSCRATCH (DUMP TYPE=SCR) are found in [Section 51](#)
- ABR Application Backup (DUMP TYPE=APPL) is in [Section 52](#).

ABR BACKUPS All the backup functions of ABR share these common characteristics:

- the backups are in standard full-volume (FDR) or data set (DSF) backup format, described in [Section 02.02](#). If necessary, data can be restored from ABR backups with FDR or FDRDSF, but ABR automates the restore process. SAR (Stand Alone Restore) can also be used to restore from ABR full-volume backups without an operating system.
- although ABR backups require a separate backup data set for each DASD volume processed, ABR will automatically stack multiple backup data sets on tape, creating multi-file tapes, to make best use of today's high-capacity tape volumes (such as IBM Magstar and StorageTek Redwood or 9840). If necessary, multiple output tape volumes are used. No special JCL is required since ABR will handle the file creation internally.
- ABR backups to disk are also supported. ABR will automatically allocate the backup data sets on the backup volumes you specify. However, disk backups are rarely used with Volume Backups.
- the output devices (tape and/or disk) are specified in the ABR batch JCL. However, you only need to identify the output device; ABR will automatically name and create every backup data set and catalog it if required.
- You can specify in the JCL that two copies of each backup are to be created, even though the input disk is read only once. This is often used to create one copy for onsite recovery and a second for offsite storage for disaster recovery. In ABR these are known as COPY 1 and COPY 2 (see [Section 50.04](#) "FDRABR Job Control Requirements" for details). The ABR utilities FDRTCOPY and FDRTSEL ([Section 60](#)) can be used to create COPY 2 and also COPY 3 through 9.
- If multiple output devices are specified in the JCL, ABR will automatically use internal subtasking to backup more than one DASD volume concurrently.

50.01 CONTINUED . . .

**ABR VOLUME
BACKUPS**

ABR Volume Backups have two purposes:

- they provide you with the ability to quickly and easily recreate entire DASD volumes which have been lost due to hardware problems or other damage. At a disaster/recovery site, you can quickly recreate the volumes necessary to run your system.
- they allow you to restore recent backups of individual data sets, specifying only the data set name.

ABR Volume Backups allow you to create daily backups of all of your DASD volumes (or a subset if you wish). However, you don't have to backup all the data every day. A daily full-volume backup of each of your DASD volumes would be too time-consuming, so users who do not have a product like ABR usually content themselves with weekend full-volume backups and leave themselves open to loss of a weeks worth of updates if a failure or disaster occurs late in the week.

ABR Volume Backups consist of:

- full-volume backups which will include all data sets which reside on each DASD volume processed. Most users do full-volume backups once a week, but you can run them on whatever schedule suits you. However, you **must** do a full-volume backup of each DASD volume involved periodically; Innovation recommends that you do this at least once a month. If you do not have time on weekends to backup all your volumes, ABR includes a facility to select certain volumes each day for full-volume backups, doing incremental backups on the rest.
- incremental backups which will backup only those data sets which have been updated since the last ABR Volume Backup (either full or incremental) was run, plus descriptive data such as the label track, VTOC, VTOCIX, and VVDS. As described later under "ABR Volume Restores", these incremental backups are sufficient to recreate an entire volume exactly as it existed when the last backup was taken, yet the elapsed time of the backup is much less than that of a full-volume backup (depending on the amount of data updated). This allows you to create daily backups of your disk volumes, so that you can do anything from restoring individual data sets to recreating entire volumes. Most users run incremental backups once a day, except on the day that the full-volume backup is scheduled. Incremental backups must be kept at least until the next ABR full-volume backup is run.

ABR volume backups do not use a data base for recording the backups. Information about backups is stored in 4 places:

- every volume to be processed for ABR backups must have a special DSCB in the VTOC of the volume, known as the "ABR Model DSCB". It is called a model because the data set has no extents and occupies no tracks on the volume, similar to the model DSCB used with GDGs. The DSCB itself is used for storing options for ABR processing of the volume and for recording data about volume backups for that volume, such as the date of the most recent backup. An ABR utility is used to create the ABR Model DSCB.
- the backup data sets created by ABR are cataloged in a normal MVS ICF catalog. The backup data sets have a fixed naming convention starting with a specified high-level index (default is FDRABR) which is usually assigned to a separate catalog known as the ABR catalog. Details of the name are described later.
- Information about the backups of individual data sets is stored in reserved fields of the Format 1 (F1) DSCB of each data set on disk.
- for data sets which have been scratched, so that their F1 DSCB is no longer available, ABR records backup information in a special MVS ICF catalog entry, called the "scratch catalog" entry, usually part of the ABR catalog.

50.01 CONTINUED . . .

**FDR
INSTANTBACKUP**

FDR InstantBackup is a separately priced member of the FDR family which enhances FDRABR which allows you to take "instant" point-in-time volume backups of online disk volumes with minimal disruption of normal processing. It requires special hardware features available in some disk subsystems. This section does contain brief descriptions of the enhancements present when ABR and FDR InstantBackup are both licensed. FDR InstantBackup is described in detail for each supported hardware platform:

- EMC Symmetrix with the Timefinder feature - described in [Section 25](#)
- IBM RVA (Ramac Virtual Array) or StorageTek Iceberg or SVA (Shared Virtual Array), with the Snapshot feature - described in [Section 26](#)
- Hitachi Data Systems 7700E with the ShadowImage feature - described in [Section 27](#)

FDR InstantBackup also enables support for the HSDM (High Speed Data Mover) feature available on some disk subsystems. HSDM allows FDR to backup and restore data in an internal compressed format. If your disk subsystem includes HSDM, you can invoke HSDM support by adding the DCT=YES operand on DUMP statements. [See Section 80.33](#) for details.

**ABR VOLUME
RESTORES**

If you have disk failures, or have to restore at a disaster site, ABR automates the restore process. It locates the most recent full-volume backup (or an earlier one if you prefer) and all the incremental backups that followed it. ABR then reads the most recently-created incremental, restoring the label track, VTOC, VTOCIX, VVDS and any data sets on that backup. It then reads back through the preceding incrementals and the full backup, restoring the most recent copy of each data set. The result is a volume which looks **exactly** like the original volume did at the time of the last incremental. All data sets are in their original locations, with the exact same allocation characteristics.

Restoring an entire DASD volume is simple; just provide the volume serial. ABR will identify and mount all of the backup tapes required. Multiple DASD volumes can be restored in one ABR job, one at a time, or multiple ABR jobs can be used to do the restores in parallel as long as the same backup tapes are not required for 2 restores.

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

**ABR DATA SET
RESTORES**

If you need to restore individual data sets, ABR tracks which full or incremental backup contains the most recent backup of each data set. Restore is as simple as providing the data set name; ABR will locate and mount the proper backup and restore the data set; the restored data set can be renamed and/or directed to a new volume, but by default it is restored to its original volume with its original name. ABR has an option to track up to 13 previous backups for every data set, allowing you to easily restore earlier versions. ABR data set restores follow the same rules as DSF restores described in [Section 20.01](#).

If more than one data set (or a multi-volume data set) is requested, ABR uses the backup information for all requested data sets to construct a list of the backup data sets which must be read. In most cases, it will dynamically allocate the restore backup data sets (on tape or disk) and read them. Tape backups are sorted; if multiple backup files on the same tape are required, they will be read in order without dismounting the tape in between.

FDRCLONE

FDRCLONE is a separately-licensed enhancement to FDR/ABR. It is described in detail starting in [Section 50.70](#).

FDRCLONE is a facility which "clones" disk volumes or data sets belonging to one MVS system to another MVS system, either an LPAR or a separate system. Its input is your normal FDRABR VOLUME backups, but data sets are actually restored only when they are needed. You can clone:

- All disk volumes in a data center
- selected disk volumes
- selected data sets

50.01 CONTINUED . . .

As cloned data sets are needed, they are dynamically restored from regular FDRABR volume backups (full-volume and incremental ABR backups) as described in this section. Only data sets which are actually needed by batch jobs or TSO users will be restored, so the total size of the restored data may be much less than the total in use at your home site.

**GENERATIONS
AND CYCLES**

In order to name the backup data sets created by ABR Volume Backups and to record those backups, ABR uses generation and cycle numbers.

The **Generation number** associated with a given DASD volume is incremented every time you take a ABR full-volume backup of that volume. The very first backup of a volume (which must be a full-volume) will be generation 1, the next full-volume backup will be generation 2, etc.

Within each generation, backups are assigned **Cycle numbers**. The full-volume backup which starts a generation is always cycle 00. Subsequent incremental backups will increment the cycle number, so the first incremental after the full-volume is cycle 01, the next is cycle 02, etc. The maximum cycle number is 63; if you don't run a full-volume backup before the cycle number reaches 63, a full-volume backup is forced in place of the next incremental backup.

The most recently used generation and cycle numbers for a given DASD volume are stored in the ABR Model DSCB on that volume, so ABR always knows the next cycle number or generation number to be used when you run a Volume Backup.

**VOLUME
INITIALIZATION**

Before a DASD volume can be processed by ABR Volume Backup, it must be initialized for ABR processing. You initialize a volume for ABR by executing an ABR utility to place an "ABR Model DSCB" in the VTOC of the volume. It is called a model because the data set has no extents and occupies no tracks on the volume, similar to the model DSCB used with GDGs.

The ABR Model DSCB can be created with the ABRINIT statement of the utility FDRABRM, described in [Section 50.40](#). It can also be created with the ABR ISPF dialog (option A.I.8) as described in [section 50.46](#); this is the preferred method.

The presence of the ABR Model DSCB is sufficient to enable the volume for ABR Volume Backups. It is used for several purposes: it contains defaults and options used by Volume Backups, such as the retention period and the number of generations to keep, and it also contains values updated during the backup, such as the latest generation and cycle numbers and the expiration date of the last full-volume backup. The defaults and options can be changed at any time by FDRABRM or ISPF option A.I.8.

The ABR Model DSCB will have a name of "FDRABR.Vvolser" where "volser" is the volume serial of the volume. "FDRABR" is the default first index; you can change it during installation of ABR but once you start running ABR backups you can no longer change it without losing all of your previous data.

On non-SMS-managed volumes, the ABR Model DSCB is an uncataloged data set, but on SMS-managed volumes it will be cataloged to meet SMS requirements. When converting ABR-initialized non-SMS volumes to SMS management, you should always use the FDRCONVT utility described in [Section 70](#) since the IBM CONVERTV function will not convert the ABR Model DSCB properly.

On SMS-managed volumes, ABR Volume Backup will also test the SMS storage group of the volume for the attributes AUTO-DUMP and AUTO-BACKUP. Volumes with AUTO-DUMP=YES but AUTO-BACKUP=NO will be processed for full-volume backups only; they are bypassed on incremental backups. If both attributes are YES, they will be processed for incremental backups as well. If both are NO, the volume will not be processed by ABR Volume Backups. If you do not wish to test the SMS storage class attributes, specify SMSCONSTRUCT=NO on the DUMP statement. SMS volumes must still be initialized with an ABR Model DSCB even if the storage group allows the backup.

NOTE: Although the term "initialization" sounds drastic, ABR volume initialization will not disturb any existing data on the volume. However, if ABR finds non-zero data in the reserved DSCB fields it uses for storing backup data, the initialization will fail; if you get a FDR420 message with text referring to "NON-ZERO" data, please call Innovation for assistance.

50.01 CONTINUED . . .

**TAPE FORMAT
AND NAMING
CONVENTIONS**

As indicated earlier, the backup files created by ABR Volume Backups are in standard FDR (full-volume) or DSF (incremental) format. In this format, each backup file can contain data only from one DASD volume. So, if multiple DASD volumes are to be processed, ABR must create multiple backup files on the output tape. Unless you restrict ABR to processing only a single disk volume in an ABR step, ABR will **always** create multiple backup tape files.

Since ABR must be able to uniquely name each tape file, and must be able to record the backup file in a way that it can easily be retrieved, ABR uses a special naming convention for the Volume Backup files. The name contains the disk volume serial, and the generation and cycle numbers, so each ABR Volume Backup from a given disk volume will have a unique name.

The format is: **abrindex.Vvvvvvv.Cnggggcc**
where:

- abrindex** is the ABR high-level index from the FDR Global Option Table. It is usually "FDRABR". It can be changed when ABR is installed but must NOT be changed once you start using ABR.
- vvvvvv** is the volume serial of the disk volume from which the data sets were dumped. ABR creates one backup file for each disk volume processed in a given ABR run.
- n** is the copy number. ABR always creates COPY 1 and can optionally create COPY 2. Additional copies up to 9 can be created by FDRTCOPY or FDRtsel (see [Section 60](#)).
- gggg** is the generation number (4 digits)
- cc** is the cycle number (2 digits)

Examples: FDRABR.VPROD01.C1023703 (Copy 1, Gen 237, incremental Cycle 3 for volume PROD01)
FDRABR.VTSOPK1.C2001200 (Copy 2, Gen 12, full-volume Cycle 00 for volume TSOPK1)

Warning: you cannot override this naming convention, except for changing the high-level index during installation. As you will see in [Section 50.04 "Volume Backup Job Control Requirements"](#) you must specify a data set name on the output (TAPEX) DD statements to satisfy MVS requirements, but that name will be ignored and the ABR-generated name used in its place.

The backup files created by ABR Volume Backup will always be cataloged. The high-level index chosen (e.g., "FDRABR") is usually assigned to an ICF catalog designated for ABR use, called the "ABR catalog". The ABR backups are cataloged like any other non-VSAM data set. If necessary, you can refer to an ABR Volume Backup directly in JCL. However, the ABR restore process makes this unnecessary, since it identifies the required backup data set names, locates their volumes in the ABR catalog and automatically mounts them.

If the backups are on high-capacity tape cartridges, such as IBM 3590 Magstar (in native mode or 3490E emulation) and StorageTek 9840 and Redwood, ABR will record a hardware "block id" in the catalog pointing to the beginning of each backup file. During restores, this block id will be used to invoke a high-speed hardware search to position directly to the beginning of the backup file. This will significantly reduce restore times.

50.01 CONTINUED . . .

**BACKUP
RETENTION
AND TAPE
MANAGEMENT**

Depending on the requirements for recovery established by your installation, you may choose to control the retention of ABR Volume Backups in several ways:

- You can let your tape management system expire backups based on a retention period or expiration date. This is known as "date control".
- You can let ABR decide when to expire backups based on the number of generations which exist for a given DASD volume. ABR has no formal interface to any tape management system, but it will uncatalog obsolete generations so you must use the "catalog control" feature of your tape management system.
- If you have no tape management system, you must establish the necessary manual procedures for identifying and expiring ABR backups. ABR will still uncatalog obsolete generations.

NOTE: if you have a tape management system from any vendor, you should enable the TMS (Tape Management System) option in the FDR Global Option Table ([see Section 90](#)). The TMS option changes slightly the way that ABR handles files on tape to be compatible with restrictions of some such systems.

When a **full-volume backup** is taken, an expiration date is assigned. It can come from several sources, in this order of priority:

- if the RETPD= operand is specified on the ABR DUMP statement, ABR uses that to calculate an expiration date which will be assigned to all COPY1 (TAPEX) backups created in that ABR step. If the RETPD2= operand is not specified but COPY2 (TAPEXX) backups are also created, the same expiration is assigned to those backups.
- If the RETPD2= operand is specified on the DUMP statement, ABR uses that retention period to calculate an expiration date which will be assigned to all COPY2 (TAPEXX) backups created in that ABR step.
- if the TAPEX or TAPEXX DD statement pointing to the backup tape contains the EXPDT= operand, the date specified is assigned to all backups written to that DD. For some tape management systems this value can also be a keyword; for example, EXPDT=99000 indicates "catalog control" to many tape management systems.
- If the DD statement contains the RETPD= operand, MVS uses that to calculate a expiration date, which ABR will assign to each backup written to that DD.
- When a volume is initialized for ABR processing, a default retention period is specified and is recorded in the ABR Model DSCB. If none of the above RETPD/EXPDT operands are given, that default retention period is used to calculate an expiration date which will be assigned to the backup of this particular volume. Separate default retentions can be specified for COPY1 (TAPEX) and COPY2 (TAPEXX) backups; if the COPY2 default is not specified, the COPY1 default is used for all backups.

By whatever means it is calculated, the expiration date of the full-volume backup is recorded in the ABR Model DSCB. If both COPY 1 and COPY 2 of a backup are being created, the expiration of each copy is recorded separately in case they are different.

The expiration date will be recorded in the tape labels of the backup file and will also be recorded by your tape management system, if you have one. If the expiration date is a real date (not a keyword such as 99000) then your tape management system will probably return the tape to scratch status on that date.

50.01 CONTINUED . . .

**BACKUP
RETENTION
AND TAPE
MANAGEMENT
(continued)**

For **incremental backups** the rules are similar. The first 4 bullets above apply. However, if none of the RETPD/EXPDT operands are given, the incremental backup is assigned the same expiration date as the most recent full-volume backup, as recorded in the ABR Model DSCB. This causes an entire generation, from the full-volume backup that began it to the last incremental, to expire on the same day. If the full-volume backup created both COPY 1 and COPY 2, and the incremental backup does the same, the expirations of the full COPY 1 are copied to the incremental COPY 1 for that volume, and likewise for COPY 2. If the full-backup only created COPY 1, its expiration is copied to all incremental backups of that volume.

If an incremental backup is requested, but ABR detects that the expiration date of the last full-volume backup in the ABR Model DSCB is past, ABR assumes that full-volume backup is no longer available. Since the full backup is required for volume restore, ABR forces a full-volume backup, starting a new generation, in place of the incremental. You must take steps to insure that your normal backup procedures will request a periodic full-volume backup before the expiration date of the last full backup is reached.

If you have a tape management system, that system will usually scratch a backup tape (or multi-volume tape set) when every backup file on the tape has reached its expiration date. Since the backups created on a tape by a given ABR run will usually have the same expiration, this is not a problem. But there are options to place backups with varying expirations on the same tape, so you should beware of creating files on the same tape whose expirations are widely separated, since this may cause the tape to be retained longer than necessary. The tape management system will usually uncatalog tape files that it expires, so ABR considers an uncataloged ABR Volume Backup to be unavailable. This technique allows your tape management system to decide when to expire ABR Volume Backups.

ABR has an additional mechanism for expiring Volume Backups. Another value that you can specify when you initialize a volume for ABR processing (stored in the ABR Model DSCB) is the number of generations to keep for the volume. Whenever a new generation is started (full-volume backup), ABR will subtract that value from the current generation number, and then will uncatalog that older generation, including all of its cycles and both COPY 1 and COPY 2. So, if you run your full-volume backups on a regular schedule, such as weekly, you can easily control retention for each disk volume by setting the number of generations to be kept. This allows ABR to determine when the backups should expire if you use the "catalog control" option of your tape management system.

For example, if a disk volume is initialized for 5 generations (a month if the full backups are weekly), and incrementals are done every weekday, the existing backups might look like:

<u>GEN</u>	<u>CYCLE</u>	
0003	00,01,02,03,04,05	
0004	00,01,02,03,04,05	
0005	00,01,02,03,04	(this week had a holiday on which no backups were run)
0006	00,01,02,03,04,05	
0007	00,01,02,03,04,05	

If a new full-volume backup is taken, it becomes generation 0008 cycle 00, and ABR will uncatalog all cycles (00 through 05) in generation 3. If you use tape management catalog control for these backups, they will automatically expire when ABR uncatalogs them.

Although the ABR backups will be uncataloged, either by ABR generation processing or by tape management expiration processing, it is possible that some generations or cycles will be left cataloged even when they are no longer required. The ABR utility FDRABRCM ([see Section 50.50](#)) includes a PURGE BACKUP function to clean up such orphans. If run with no other parameters, it will examine every Volume Backup recorded in the ABR catalog and compare it to the generations limit recorded in the ABR Model DSCB on that disk volume; backups in generations older than that are uncataloged. You may want to run PURGE BACKUP monthly to be sure no orphan backups are still cataloged.

50.01 CONTINUED . . .

**BACKUP
RETENTION
AND TAPE
MANAGEMENT
(continued)**

You can use a combination of these retention techniques. This is usually used to manage the incremental backups differently from the full backups. For example, assume you do weekly full backups and daily incrementals and want to keep the backups for the last 6 weeks. However, since this could occupy many tape volumes, you only want to keep the daily incremental backups for the most recent 2 weeks, keeping only the full backups for the remaining 4 weeks. To accomplish this:

- set the number of generations in the ABR Model DSCB to 6
- execute the full-volume backups with catalog control (EXPDT=99000 for many tape management systems)
- execute the incremental backups with RETPD=14 (2 weeks)

Your tape management system will automatically expire and uncatalog the daily incremental backups after 2 weeks, but the full backups will be kept until they are uncataloged by ABR (6 weeks). It is also possible to put all the backups under catalog control but accomplish the same thing using the PURGE BACKUP function of FDRABRCM ([See Section 50.50](#)).

If you create both COPY 1 and COPY 2 of your backups, you can use one copy for onsite recovery (for damaged disks or restore of individual data sets) and the other copy for offsite recovery at a disaster/recovery site. Most tape management systems include vaulting support, allowing you to select tape to be sent offsite and returned when no longer required. Most users send COPY 2 offsite. Since the data set name used by ABR Volume Backup includes the copy number (described earlier), you should be able to easily send COPY 2 backups offsite while retaining COPY 1 onsite.

50.01 CONTINUED . . .

**VOLUME
BACKUP
EXECUTION**

To execute ABR and perform Volume Backups, you must create ABR jobstreams and execute them at appropriate times. [Section 50.20](#) contains examples of such jobstreams which you can customize for your installation; all of the examples shown in this manual are also in the JCL library loaded as part of ABR's installation ([Section 90](#)). If you have an automation or scheduling product in your installation, you may want to use it to schedule daily ABR Volume Backups at appropriate times.

Since ABR requires periodic full-volume backups, you will want to schedule TYPE=FDR backups at regular intervals, such as every weekend or every other weekend, running TYPE=ABR incremental backups on other days. Alternately, you can use TYPE=AUTO to allow ABR to decide when to do incremental and when to do full-volume backups; this can be used to spread your full-volume backups out over the week.

NOTE: Even though you are executing an incremental backup (TYPE=ABR, AUTO, or DSF), ABR may force a full-volume backup to be taken of certain disk volumes if:

- 1) it is the first ABR Volume Backup of the volume.
- 2) the last cycle number was 63, which is the maximum cycle allowed.
- 3) the expiration date of the latest full-volume backup, as recorded in the ABR Model DSCB, is past.

The expiration date test can be overridden by the DATEP=NONE operand on the DUMP statement.

For full-volume recovery, you will also have to create ABR jobstreams; examples are in [Section 50.23](#) and the JCL library. Since full-volume recovery at your home site is unusual, we suggest that you have a model RESTORE TYPE=FDR jobstream which can be modified as necessary. For use at a disaster/recovery site, we suggest that you create all the jobstreams necessary to recover volumes from Volume Backups and send them to offsite storage so that you do not have to create them under pressure at the site.

For data set recovery from Volume Backups, there are several alternatives:

- you can permit users to restore data sets directly by submitting their own RESTORE TYPE=ABR job. This may require mounting backup tapes for each data set required. If FDR security is enabled, users will need UPDATE or ALTER authority to the data sets.
- you can allow users to add requests to a remote queue (see details in [Section 50.03](#)). In this case Operations (or your automation software) must submit a RESTORE TYPE=ABR job at regular intervals (e.g., every hour, every 4 hours) to process any requests on the queue. If multiple data sets are being restored, this minimizes tape mounts. If FDR security is enabled, users will need UPDATE or ALTER authority to the data sets.
- You can require that users communicate their data set restore requests directly to someone in the Data Center, who will submit the required RESTORE TYPE=ABR job.

[Sections 50.21](#) and [50.22](#) and the JCL library contain examples of data set restore jobstreams. For the first 2 bullets, the ABR ISPF panels support direct restore and remote queue restore, so that users do not need to know anything about ABR JCL when restoring their data sets. [See Section 50.30](#) for details.

50.02 DATA SET SELECTION

FULL VOLUME BACKUP Naturally, during a full-volume ABR backup all data sets on the volume are always backed up. There are no options to select or exclude particular data sets. The backup also includes the label track, VTOC, VTOCIX and VVDS (if the latter 2 exist).

INCREMENTAL BACKUP During an incremental backup (DUMP TYPE=ABR or TYPE=AUTO) data sets will be selected for backup from a given disk volume by the following rules:

- **LABEL TRACK** (cylinder 0 track 0) - always backed up
- **VTOC** - always backed up.
- **VTOC INDEX (VTOCIX)** - always backed up if it exists.
- **VVDS** - always backed up if it exists.
- **PROTECT LIST** - data sets named in the ABR Backup Protect List will not be included in any incremental backup. The Backup Protect List is defined and enabled by the ABR ISPF Install Dialog described in [Section 90](#).
- **SMS MANAGEMENT** - if SMSMANAGE=YES is specified on the ABR DUMP statement, the management class of every SMS-managed data set is checked. If the management class specifies that the data set is not eligible for "auto-backup", it will be automatically excluded.
- **SELECT and EXCLUDE STATEMENTS** - ABR will backup or exclude any data set selected by the user. SELECT and EXCLUDE control statements will override all backup selection rules except for the backup of the VTOC and the PROTECT list. If UPDATE is specified on a SELECT statement the data set will only be dumped if the update (changed) flag is set in its Format 1 DSCB.

Warning: if you exclude data sets from incremental backups with the PROTECT LIST, SMS management class or EXCLUDE statement and they were updated since the last full-volume backup, a full-volume recovery will not properly restore the contents of those data sets.

- **TEMPORARY DATA SETS** - ABR will not backup temporary data sets. A temporary data set is defined as a data set whose name starts with the characters 'SYS', and contains a '.' in position 9 and 10, and a period in position 17.
- **UPDATED DATA SETS** - unless excluded by the Protect List, SMS or EXCLUDE statements, ABR will backup all data sets on a disk volume which have been updated since the last backup. ABR examines the update (changed) flag in the DS1DSIND field of each Format 1 DSCB. If the data set has been updated, it will be backed up and the update bit will be reset so that it will not be backed up again on the next incremental unless it is updated again. This is the usual way that data sets are selected for incremental backups.
- **UPDATED ICF VSAM** - IBM only sets the update flag in the DSCB of the base data component, and only on the first volume in the case of multi-volume clusters. Whenever ABR selects any component of a ICF VSAM cluster from a volume, it will automatically select all components of that cluster (including alternate indexes) on the same volume. Multi-volume data components, index components on a volume without the data component and key ranges (all volumes) will not have the update indicator set. ABR will always dump these components if found on the volume, except for an alternate index residing on a volume by itself.
- **NEW DATA SETS** - ABR will also backup a data set if it does not have a current ABR backup. This implies that it was created since the last ABR Volume Backup of this volume.
- **CATALOGS** - all ICF VSAM and OS CVOL catalogs on the volume are backed up.
- **REMOTE QUEUE** - ABR supports user requests for data sets to be included in incremental backups even though they have not been updated. The ABR ISPF panels support this function. Depending on the installation options, when a user requests Remote Queue backup of a data set, it will either set the update flag in the Format 1 DSCB of the data set, or it will add a statement to a Remote Queue data set requesting the backup. In the latter case, options in the ABR backup step control whether these requests will be honored.

50.02 CONTINUED . . .

DATA SET BACKUP ABR Volume Backup includes an option to only backup specific data sets. This is requested by TYPE=DSF on the ABR DUMP statement. In this mode, SELECT and optional EXCLUDE statements specify the data sets to be included in the backup. No other data sets (not even the VTOC) will be included. However, you should be aware that if you include TYPE=DSF runs with regular incremental backups, ABR full-volume recovery may not work correctly. But you will be able to restore individual data sets from any type of volume backup.

FULL-VOLUME RECOVERY ABR Full-volume recovery from Volume Backups does not recover individual data sets. Instead it reads the most recent incremental backup to recover the most recent copy of the label track, VTOC, VTOCIX and VVDS, so that all data sets appear to be at their most recent locations. Then it recovers the most recent copy of every allocated data track, starting with the most recent incremental backup and working backwards to the full-volume backup that started the generation. So, if a data set was backed up on several different incremental backups, only the data from the most recent backup is restored.

To determine the backups that must be read, ABR does this:

- If the volume being recovered still exists and is readable, it will read the ABR Model DSCB from that volume. That DSCB contains the generation and cycle numbers of the most recent backup for that volume. The recovery will start with that last generation/cycle and work backwards to cycle 00 (full backup).
- If the volume is not accessible or the ABR Model DSCB does not exist, ABR will search the ABR catalog for the most recently created backup for this volume, based on the date of the catalog entry. The recovery will start with the generation/cycle of that backup and work backwards to cycle 00 (full backup). Since you usually do full-volume recovery after a hardware failure, or at a disaster site, this is the most frequently used technique.
- If you want to recover starting with a cycle other than the most recent, or recover from a generation other than the current generation, you can specify the generation and/or cycle number to be recovered.

WARNING: *if you use the DEFrag function (part of IBM's DFSMSdss and DFDSS products) to consolidate free space on volumes for which you do incremental backups, the ABR Full-Volume Recovery will not work properly. ABR is dependant on the location of data sets to properly reconstruct the volume, so when a data set's location is moved, it must be backed up again on the next incremental backup to enable proper full-volume recovery. DEFrag does not turn on the UPDATE flag in data sets that it moves, so ABR doesn't know that they need to be included in the next incremental. If you must use DEFrag, you should run it immediately before the ABR full-volume backup (TYPE=FDR) for the volume. A better choice would be to use COMPAKTOR (Section 40) which does set the UPDATE flag for all data sets moved.*

50.02 CONTINUED . . .

**DATA SET
RESTORE**

Since the ABR Volume Backups are standard FDR backups, you can restore any individual data set from these backups. For recently created backups, ABR can automate the entire restore process, so that you do not have to have any knowledge of which generation and cycle contains the backup that you want restored; you usually don't even have to know what disk volume it was on.

As described in [Section 50.01](#), ABR Volume Backup records the most recent backup of each data set in fields in the Format 1 DSCB of that data set. Actually it only records the cycle number that contains that backup; since that cycle must be part of the current generation, the current generation number is obtained from the ABR Model DSCB on the volume and the disk volume serial completes the information necessary to construct the name of the ABR backup file on tape. That name is then located in the ABR catalog and the required tape is mounted.

For volumes which have been initialized with the OLDBACKUP option in the ABR Model DSCB (see "Volume Initialization" in [Section 50.01](#)), ABR will also record up to 13 previous backups of a given data set. Every time a data set is backed up by Volume Backup, the last "current backup" is moved into the most recent "old backup" slot and the rest of the old backups shifted down, losing the oldest. If you need to restore a backup other than the most recent, you can either request it by relative backup number (0=current, 1=next oldest, etc.) or display the recorded backups under ISPF and select the version desired. For volumes where OLDBACKUP is not enabled, ABR can only restore the most recent backup automatically.

Since you may wish to restore the backup of a data set which has been scratched (deleted) from disk, making its Format 1 DSCB unavailable, ABR has an optional facility called the SCRATCH catalog. Details of enabling the DADSM Preprocessing Exit, which implements the SCRATCH catalog, are found in [Section 90](#). When the exit is enabled, and data set which has a current backup recorded in its Format 1 DSCB is scratched or renamed, the exit will extract the backup information (the current backup plus up to 3 previous backups if present) and record that information as a catalog entry in the SCRATCH catalog (which is usually the same as the ABR catalog).

To automatically restore a data set from a Volume Backup, ABR must find the Format 1 DSCB (VTOC entry) of that data set, since backup information is kept in the F1 DSCB. First, ABR must determine the disk volume serial from which the data set was dumped. If you do not provide the volume serial, ABR will search the catalog for the data set and get the volume serial from the catalog entry. ABR then reads the F1 DSCB of the data set from the VTOC of that volume to determine the generation and cycle of the required backup, either the current backup or an older backup.

If the data set was not found in the catalog or in the VTOC, or it has no current backup recorded or the volume was offline, ABR will search the SCRATCH catalog if enabled. If you specified a volume serial, the search will be successful only if it is recorded as being scratched or renamed on that volume serial; otherwise it will select it from the SCRATCH catalog no matter what volume it was scratched from.

If you know the particular volume serial and backup generation and cycle numbers which contain the backup of the data set you want, you can specify this information via the **VOL=**, **GEN=**, and **CYCLE=** operands on the SELECT statement. ABR will still locate that backup in the ABR catalog and mount the required tape.

50.02 CONTINUED . . .

**DATA SET
RESTORE
(continued)**

When restoring a data set from Volume Backups, the original data set often still exists on disk. In that case, by default, ABR will simply overlay the existing data set with its earlier contents. But if it is not currently on disk, or if you force the restore to a new volume or a new data set name, the restore process is identical to that of FDRDSF. Please review [Section 20.01](#) for details on allocating, restoring, and cataloging data sets during the restore. [Section 50.08](#) contains details on output volume selection for ABR restores.

The data set restore process above works only when you are requesting specific individual data sets. If your SELECT statement specifies a mask (e.g., `SELECT DSN=PAYROLL.*`) or `SELECT ALLDSN`, ABR will require that you identify one particular backup data set (either a full-volume backup or an incremental backup, from one particular disk volume) to scan for the data sets. In this case ABR will not locate the F1 DSCB, nor will it check the SCRATCH catalog; only that one backup file will be read. You can identify the backup by VOL=, GEN= and CYCLE= as mentioned above

However, if the data sets to be restored are still cataloged, you can use a catalog mask (e.g., `SELECT CATDSN=PAYROLL.*`). CATDSN scans the system catalogs for data sets matching the mask and internally generates a `SELECT DSN=dsnname, VOL=vvvvvv` for each, so it will locate their backups in the F1 DSCB or SCRATCH catalog.

The **ABR Restore Protect List**, documented in [Section 90](#), can list data sets or groups of data sets which are not to be restored, similar to EXCLUDE statements. Any attempt to restore those data sets from Volume Backups or Archive Backups will fail (except for full-volume recovery from Volume Backups).

[See Section 20.02](#) for details on allocating and cataloging restored data sets.

50.03 PROCESSING OPTIONS AND REQUIREMENTS

ABR OPERATIONS The first control statement in the ABR input must be a DUMP or RESTORE statement. It may be followed by SELECT and optionally EXCLUDE statements to specify data sets which must be backed up. It may also be followed by MOUNT statements to specify the DASD volumes or SMS storage groups to be processed in this execution of ABR (DISKxxxx DD statements can also be used to specify volumes as described in [Section 50.04](#)). ABR will accept up to 100 SELECT/EXCLUDE/MOUNT statements in a single execution, unless that limit is overridden.

ABR BACKUP OPTIONS The DUMP statement contains a TYPE= operand to specify the type of backup. In each case, volumes to be processed in the ABR step are identified by MOUNT statements and/or DISKxxxx DD statements. These 4 values will invoke ABR Volume Backups:

A. TYPE=FDR

A full-volume ABR Volume Backup will be taken for every selected volume. SELECT/EXCLUDE statements are ignored; all data sets on the volumes are backed up. However, a EXCLUDE ALLDSN with VOL= or VOLG= will still exclude those volumes from ABR processing.

B. TYPE=ABR

An incremental ABR Volume Backup will be taken for every selected volume. All data sets which have been updated since the last Volume Backup of the volume, plus all newly created data sets and the label track, VTOC, VTOCIX and VVDS are backed up. SELECT statements may be used to include other data sets even if they don't otherwise qualify. EXCLUDE statements, the ABR PROTECT LIST and SMS management class may be used to exclude certain data sets. See [Section 50.02](#) for more details on data set selection.

C. TYPE=AUTO

Operates the same as TYPE=ABR **except** that it will automatically do a full-volume backup instead of an incremental at intervals. When you initialize the volume for ABR processing, one of the values you can set in the ABR Model DSCB is a cycle limit. That limit is used only with TYPE=AUTO backups. Every time an incremental backup is executed, a cycle count in the ABR Model DSCB is incremented. If the cycle count exceeds the cycle limit, the full-volume backup is forced. This can be used in place of explicit TYPE=FDR full-volume backups. It is possible to arrange TYPE=AUTO backups so that some percentage of your DASD volumes will take full-volume backups each day, with the rest doing incrementals.

50.03 CONTINUED . . .

**ABR BACKUP
OPTIONS
(continued)****D. TYPE=DSF**

Operates the same as TYPE=ABR, except that **only** the data sets identified by SELECT/EXCLUDE statements (which are required) will be backed up. The label track, VTOC, VTOCIX and VVDS will not be included in the backup. TYPE=DSF is a way of backing up only selected data sets in the Volume Backup system, but there is an important consideration: if you need to do an ABR full-volume recovery (RESTORE TYPE=FDR), and the most recent incremental is a TYPE=DSF backup it will be bypassed. So TYPE=DSF backups are not recommended; if you need to backup specific data sets, see "ABR Application Backup" in [Section 52](#).

In all these options, ABR builds a list of the DASD volumes to be processed. If the ABR step JCL contains a single TAPEX DD statement, ABR will select the first disk volume on its list and do the requested backup, creating file 1 on that output tape. When complete, it selects the next volume and creates file 2 on the tape, etc., until all volumes are processed.

However, you may have up to 9 TAPEX DDs in the step. If you have more than one, ABR will select the first "n" volumes on its list and assign them to the "n" tape drives, executing the backups in parallel with internal subtasking, and creating file 1 on each output tape. As each backup completes, ABR selects the next disk volume and assigns it to the tape which just became idle, creating the next file number on that tape. Use of this feature usually reduces the overall execution time of the backup. See "Memory Requirements" below for limitations on the use of multiple TAPEX DDs.

ABR's list of disk volumes to be processed comes from several sources: if DISKxxxx DD statements are present in the ABR step ([see Section 50.04](#)) the named volumes are added in the order they appear. If MOUNT statements are present in the input, the named volumes are added in the order of their UCB addresses. If ONLVOL is specified on the DUMP statement, any volumes named on SELECT statements are included. If ONLINE is specified on the DUMP statement, all online volumes are included.

**ABR RESTORE
OPTIONS**

The RESTORE statement contains a TYPE= operand to specify the type of backup. These 2 values will invoke restores from ABR Volume Backups:

A. TYPE=FDR

A full-volume ABR volume recovery will be performed for every selected volume. SELECT statements are used to identify the volumes to be recovered. The recovery will normally begin with the most recent incremental backup of each volume and will read back through the incrementals in the current generation, reading the full-volume backup that began the generation last. As a result, the recovered volume will be an **exact** copy of the volume at the time of the last incremental backup. If recovery of multiple volumes is requested in an ABR step, they are recovered serially (one at a time); submit multiple ABR jobs if you want to recover volumes in parallel, but beware of contention for the same backup tapes by those ABR jobs.

B. TYPE=ABR

One or more individual data sets are restored from Volume Backups. ABR will automatically identify the backup containing each requested data set and mount it. [See Section 50.02](#) for more details on data set selection and backup location.

50.03 CONTINUED . . .

**MEMORY
REQUIRE-
MENTS**

For backups specifying RTC=YES or DCT=YES, ABR requires:

- About 500K of storage below the 16MB line for programs and control blocks.
- About 2MB (2048K) of above-the-line storage for each concurrent backup.
- If COMPRESS= is specified, about 100K of below-the-line storage for each concurrent backup.

For backups specifying neither RTC=YES nor DCT=YES, ABR requires

- About 500K of storage below the 16MB line for programs and control blocks.
- About 1MB (1024K) of below-the-16MB-line storage for each concurrent backup
- If COMPRESS= is specified, an additional 1M of below-the-line storage for each concurrent backup (total 2MB per backup).

ABR will execute concurrent backups if multiple TAPEX DD statements have been provided, up to 9.

For example:

1 backup (or serial backups) without RTC/DTC – 1500K below the line

3 concurrent backups without RTC/DTC – 3500K below the line

3 concurrent backups with RTC or DCT – 500K below the line, 6MB above the line.

ABR always uses the exact memory it requires for a given function, no matter what REGION= value is specified. So, if the region is too small, ABR will fail (or some of the backups will fail); if it is too large, ABR will not use the excess so a large region has no negative impact. For this reason, you may want to specify REGION=0M in ABR JCL to request the largest available region. Alternately, you can specify REGION=xxM (xx greater than 16) to get the largest available below-the-line region and the specified amount of above-the-line region; the default above-the-line region is 32M unless your installation has changed it.

ABR full-volume recovery requires about 300MB of above-the-line storage.

ABR data set restores require a below-the-line region of 512K plus about 512 bytes for each data set or track range to be processed. ICF VSAM clusters may add an additional 1K bytes per component processed.

Some logical RESTORE operations may require additional memory, so a region of 1024K or more is recommended. You may want to specify REGION=0M to get the largest possible region. DSF will use only the storage it needs regardless of the REGION size.

**COMPRESS
OPTION**

FDRABR can be instructed to compress the data on the sequential backup file using Innovation's own proprietary software compression algorithm.

It is **not recommended** for backups to tape attached by ESCON or FICON channels because of the speed of the channel. All ESCON/FICON-attached tape drives include hardware compression which will provide savings similar to that of FDR compression. Software compression will be ignored for backups created with the HSDM disk hardware option (DCT=YES) since the data is pre-compressed.

NOTE: all FDR restores will automatically recognize a compressed backup file and decompress it. No special option is required to restore a compressed backup.

50.03 CONTINUED . . .

- DUPLICATE TAPE OPTION** ABR has an option to create a duplicate or second copy of the backup tape during dump processing. The primary copy is called COPY 1 and the duplicate is called COPY 2. The copy number becomes part of the name of the backup data set as described in [Section 50.01](#).
- While dumping a disk to a TAPEX DD statement, the duplicate backup will be written to the TAPEXX DD statement (same "x" value twice) if it is present. You may have TAPEXX DDs for some TAPEX DDs and not for others in the same step but this is not recommended with ABR since you cannot predict which disks will be written to which tapes.
- Memory requirements do not increase with the use of the duplicate tape option.
- STORAGETEK EXHPDM SUPPORT** FDRABR Volume Backups support the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM takes multiple concurrent tape outputs (such as FDRABR backup TAPEX or TAPEXX DD statements) and directs them to a smaller number of tape drives, interleaving the data in a single tape file. ExHPDM is invoked by adding the SUBSYS= operand to the TAPE DD statements. [See Section 80.33](#) for more details.

50.03 CONTINUED . . .

SECURITY Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system which supports the SAF (Security Authorization Facility) interface, such as ACF2 or TOP SECRET, you can enable the ALLCALL option. For ABR Volume Backup this results in these security checks:

- for full-volume backups (DUMP TYPE=FDR plus those forced under TYPE=AUTO or ABR), ABR will check to see if your userid has at least READ authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority or the input volume is not protected by your security system, then ABR will check if you have at least READ authority under the DATASET security class to every data set on the input volume. If you don't have volume authority or authority to every data set on the volume, the backup is terminated.
- for incremental backups (DUMP TYPE=ABR, AUTO or DSF), ABR will always check to see if your userid has at least READ authority to the entire input volume as defined above. If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority, then ABR will check if you have at least READ authority under the DATASET security class to every data set being backed up. Any data sets to which you are not authorized will be bypassed with an error message; this may impact the ability to do a proper ABR full-volume recovery.
- for full-volume recoveries (RESTORE TYPE=FDR), ABR will check to see if your userid has UPDATE authority to the entire output volume. If not, the restore is terminated. But if the output volume is not protected by your security system, the restore will be done with no additional security checks. For this reason, Innovation recommends defining volume-level security rules to control full-volume restores.
- for data set restores (RESTORE TYPE=ABR), ABR will check if you have at least UPDATE authority under the DATASET security class to every data set restored. Any data sets to which you are not authorized will be bypassed with an error message. If an output data set must be allocated, the operating system will check if you have CREATE/ALLOCATE authority for the data set (this is done even if ALLCALL is not enabled).

Innovation recommends that volume security rules be established for the userids which will be executing ABR Volume Backups and full-volume recoveries. This will minimize the security checking overhead as well as providing clearcut control over who can execute the backups. Under RACF, volume security is controlled by the DASDVOL security class, and other security systems have similar facilities (see their documentation). Under RACF, it is also possible to execute under a userid with OPERATIONS authority, which automatically grants DASDVOL and DATASET authority.

50.03 CONTINUED . . .

**DATA SET
ENQUEUE
OPTION**

You can request, via the DSNENQ= operand, that each data set being dumped or restored be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name of SYSDSN and a minor name of the data set name itself, for each selected data set found in the VTOC of the input disk; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. Note that FDR cannot tell if the data set is being used for input or output. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations. If you have requested data set ENQs, any data set that is in use will cause a FDR158 warning message to be printed; this will set the job error flag and will cause a return code 12 when the step is complete (see "Step Termination" below). If you don't want in-use data sets to be considered an error, specify the ENQERR=NO operand; this prints the FDR158 message without setting the error flag.

Optionally you can request that inactive data sets be enqueued to ABR during the backup, to insure that no other job or TSO user can access the data set until the backup is done.

For backups, in-use data sets will still be dumped by default, but you must be aware that the backups of data sets which are being updated during the backup may be unusable, depending on the nature and format of the data. If you wish to bypass the backup of active data sets during an incremental backup, specify the ENQERR=BYPASS operand.

For restores, ABR will attempt to enqueue any data sets that it allocates on the output disks, to insure that no other task tries to use them until the restore is complete, but if the ENQ fails, the data set is still restored. But for existing data sets, if the ENQ fails, the restore will be bypassed. Data sets are not enqueued during full-volume recovery.

The DSNENQ= operand has 4 possible values:

USE – data sets will be enqueued for the duration of the backup or restore on this disk volume. For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option. **This is the default for TYPE=ABR restores.**

TEST – data sets will only be tested to see if they are enqueued to another task at the time that the backup or restore on this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the backup or restore is proceeding.

HAVE – The data sets will be enqueued for the duration of the backup or restore. If a data set is in use, the MVS operator must interact with ABR to decide how to proceed; a message (FDRW27) is issued to the MVS console, and the operator can respond:

WAIT - wait for the data set to become available; if it is not eventually dequeued, the ABR job may time out, so the operator must know which data sets are in use by long-running jobs or tasks.

NOWAIT - do not enqueue the data set. The FDR158 warning message is issued.

RETRY - try the enqueue again. If it fails again, the FDRW27 message is reissued.

NONE – No data set ENQ will be issued. **This is the default for backups.**

NOTE: If a data set name appears in a DD statement with DISP=SHR within the ABR job (not necessarily in the ABR step), and you specify DSNENQ=USE, HAVE or TEST, ABR will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the ABR job ends.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Use DSNENQ= to prevent other tasks from updating (or reading) data sets being dumped or restored. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

If HFS=QUIESCE is specified, special processing is done for HFS data sets (Hierarchical File System, used by OS/390 Unix System Services). If the SYSDSN ENQ cannot be acquired, this may mean that the file system is mounted to OS/390 Unix, so FDR will attempt to quiesce the file

50.03 CONTINUED . . .

system during the backup. Details on the quiesce function are found in [Section 80.11](#).

**VTOC
ENQUEUE
OPTION**

ABR also supports, via the ENQ= operand, an ENQ on the VTOC of every volume being dumped. For shared DASD, it can also invoke a hardware RESERVE on the volume during the FDR operation.

The VTOC is protected by an ENQ with major name SYSVTOC and a minor name of the volume serial. This ENQ is held by any task doing updates to the VTOC, including allocation of new data sets, extension of data sets to new extents, and scratching of existing data sets. This ENQ is normally of short duration, just for the few seconds necessary to update the VTOC, so if the ENQ is currently held by another task, ABR will wait for it to be released.

The SYSVTOC ENQ does **not** prevent access to existing data sets on the volume; it only insures that no other task is updating the VTOC while ABR is processing it. VTOC changes during a backup could result in an invalid backup.

For disks shared with another MVS system or LPAR, ENQ=RESERVE requests that, in addition to the ENQ described above, a hardware RESERVE is done on the volume. RESERVE will prevent any system from doing any I/O on the volume, except for the system that ABR is running on where only the ENQ protection applies. If you have a cross-CPU ENQ facility, such as GRS or MIM, you may be able to convert the RESERVE into a cross-CPU SYSVTOC ENQ and allow access to the volume during the operation (look up SYSVTOC in the documentation for your product).

Use ENQ= to prevent other tasks from making changes to the VTOC during the backup. Since ABR must update the VTOC at the end of the backup, to record the backup just taken, and so must protect against unexpected changes in the VTOC, ABR Volume Backups default to ENQ=ON (ENQ on same CPU). If you are executing ABR on shared DASD, specify ENQ=RESERVE. Member ENQ in the FDR ICL (Installation Control Library) has more information on VTOC ENQs

**STEP
TERMINATION**

If no errors occur during the execution of ABR, the ABR jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

- Some errors are critical. The jobstep ends immediately with a user ABEND.
- Some errors are critical only to a particular operation. For example, during a backup, some errors cause the backup of a particular disk to terminate immediately, but ABR may continue and attempt to backup other disks requested in the same step.
- Some errors are non-critical and the messages are warnings only. ABR will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with return code 12 to call your attention to the errors. Remember that RC=12 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer to get a different return code or a U0888 abend on a non-terminating error, the ABRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice or ABEND ([see Section 90](#)).

50.03 CONTINUED . . .

REMOTE QUEUE ABR includes a facility, called "remote queue" where users can request that certain data sets be included in the next incremental backup, or that certain data sets be restored from Volume Backups. Requests can be added to the queue by TSO users using the ABR ISPF dialog; there is also a batch utility for adding requests (FDRABRUT, [see Section 51.40](#))

For backups, depending on installation options, a backup request to the remote queue can simply turn on the update flag in the F1 DSCB of the data set, or it may add a statement to a remote queue data set requesting the backup. For restores, a remote queue data set is always updated. You will find more information on the setup of the remote queue data sets in [Section 90](#) (Installation). They are honored only if you setup the data sets properly and include the proper DD statements in ABR jobs to process them.

If a remote queue data set is used for backups, a special DD statement must be included in the backup JCL ([See Section 50.04](#)) pointing to the data set. ABR will add the requests (which act like SELECT statements) to the ABR control statement input.

If a remote queue is used for restores, you will probably want to run a special ABR data set restore job at intervals (perhaps several times a day) just to process the remote requests so that the users do not have to wait an excessive time for their data sets to be restored. A special DD statement must be included in the restore JCL ([See Section 50.07](#)) pointing to the data set. ABR will add the requests (which act like SELECT statements) to the ABR control statement input, which normally consists only of a RESTORE TYPE=ABR statement.

DYNAMIC ALLOCATION ABR will dynamically allocate disk volumes as needed. Except for full-volume recovery (RESTORE TYPE=FDR), there is no need to provide DISKxxxx DD statements for volumes to be processed by ABR. As long as the required volumes are online, ABR can dump from and restore to any required disk volumes.

For restore operations, ABR will dynamically allocate each required backup tape if the DYNTAPE option is specified. If required backups are on mixed device types (such as 3480s and 3490Es), ABR will automatically mount each tape on the proper device type. For automated tape libraries, a drive in the proper library will be allocated. If you have multiple tape libraries, you may need to enable the DYNDEALC option in the FDR Global Option Table ([see Section 90](#)).

For data set restores, ABR will sort the list of backup files required and mount the backup tapes in an order which minimizes the amount of tape movement required.

50.04 VOLUME BACKUP JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform dump functions:

STEPLIB or JOB LIB DD STATEMENT	If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.
EXEC STATEMENT	Specifies the program name (PGM=FDRABR), region requirement (REGION=, see Section 50.03), and optional PARM= operand. If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid DUMP statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example, <pre>//FDR EXEC PGM=FDRABR , PARM='DUMP TYPE=FDR , DSNENQ=USE ' //FDR EXEC PGM=FDRABR , PARM='DUMP TYPE=DSF / SELECT DSN=A . B . C '</pre>
	If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.
SYSPRINT DD STATEMENT	Specifies the primary output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.
SYSPRINx DD STATEMENT	Specifies the secondary output data set for messages related to the matching TAPEX DD statement. A SYSPRINx is required for each TAPEX present in the step. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.
ABRMAP DD STATEMENT	Specifies the output report data set if SIM or PRINT=ABR is specified. It is optional; if not specified the SIM report will be printed on the SYSPRINx DD statement. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.
FDRSUMM DD STATEMENT	(Optional) if present, ABR will write one-line messages for each volume dumped, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set. FDRSUMM is used only if RTC=YES or DCT=YES is specified on the DUMP statement.
SYSUDUMP DD STATEMENT	Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced: <pre>//ABNLIGNR DD DUMMY</pre>

50.04 CONTINUED . . .

**DISKxxxx DD
STATEMENT(s)**

Optionally specifies disk volumes to be processed by ABR. The format will be:

```
//DISKxxxx DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"xxxx" may be 1-4 alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) characters, "unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). You may use either DISP=OLD or DISP=SHR; it makes no difference.

The DISKxxxx DD may also specify multiple volume serials, up to 255, in the format:

```
//DISKxxxx DD UNIT=unitname,DISP=OLD,
// VOL=SER=(volser1,volser2,...)
```

All the volumes must be of the same type (e.g., 3390) and have the same capacity (e.g., 3390-3).

ABR will build a list of the volume serials from all of the DISKxxxx DDs found in the ABR step JCL; these volumes will be processed in the order that the DISKxxxx DDs are found. A maximum of 256 disk volumes can be processed in one execution of ABR; if more disk volumes are required, MAXDD= must be specified on the DUMP statement.

However, DISKxxxx DD statements are not required and there are easier alternatives. If the ONLINE operand is specified on the DUMP statement, all online disk volumes will be processed automatically. If MOUNT statements are included in the ABR input, ABR will automatically process the indicated volumes (by volser, volume serial group, or SMS storage class). If the ONLVOL operand is specified, any volumes specified by SELECT statements (from the VOL= or VOLG= operands, or selected from the catalog by the CATDSN= operand) will be automatically processed. These volumes will be added to the list of volumes specified by DISKxxxx DDs (if any). Both ONLINE, ONLVOL and MOUNT will select volumes in the order of their device addresses (UCB address order), so you may still want to use DISKxxxx DDs to specify the order of processing for certain volumes.

NOTE: If ABR encounters multiple DISKxxxx DD statements pointing to the same disk volume serial, ABR will process the volume once using the first DDNAME encountered. DISKONLx is a reserved DDNAME, used by ABR for allocation of volumes that do not have DISKxxxx DDs.

If you are using FDR InstantBackup with ABR to backup offline point-in-time copies of online disk volumes, the DISKx DD statements will allocate the online volumes. However, the PPRC=, SNAP= or BCV= operands on the DUMP statement instruct ABR to access the offline copy instead.

50.04 CONTINUED . . .

TAPEX DD STATEMENT Specifies the output device on which backup files are to be created. "x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character. A maximum of 9 such TAPEX DDs may be present in the ABR step JCL. ABR will start an internal dump subtask for each TAPEX DD present, each subtask processing one disk volume from ABR's list of disk volumes. If only one TAPEX is provided, ABR will process disk volumes one at a time; if multiple TAPEX DDs are present, ABR will process up to 10 disks in parallel (however, region requirements may limit the number of concurrent subtasks; [see Section 50.02](#)). TAPEX DDs that specify DUMMY or DSN=NULLFILE will be ignored except for simulation.

Tape Output:

If outputting to tape or cartridge, the TAPEX DD must specify:

DSN=	a data set name is required by MVS, but it will be overridden by ABR at OPEN time, so any non-temporary name is acceptable. The name you specify will not be used by ABR, but MVS will do an exclusive ENQ on this name at job initiation so each ABR job should use unique names.
UNIT=	specify a generic (e.g., 3490) or esoteric (e.g., CART) name to allocate the type of tape drive desired. If you have sufficient tape drives available, specifying a unit count of 2 (e.g., UNIT=(3480 , 2)) may reduce elapsed time (most significant for 3480 cartridge drives).
VOL=	specify a volume count, e.g., VOL=(, , 255), to prevent ABR from abending if more than 5 tape volumes are required. If no volume serials are specified, ABR will call for scratch tapes; this is recommended; however, you may specify up to 255 tape volume serials.
LABEL=	you may want to specify RETPD= or EXPTD= to identify the expiration date of the backups (retention periods may also be specified in the ABR model DSCB on each disk or on the DUMP statement). ABR does not make use of this expiration (except that the expiration dates of full-volume backups are recorded), but if you have a tape management system it will honor the expiration specified. See Section 50.01 for details on retention of ABR Volume Backups.
DISP=(NEW,KEEP)	is required; do not specify CATLG since ABR handles cataloging of output files internally.
FREE=CLOSE	do not specify FREE=CLOSE since it will cause ABR to fail when it tries to create a second file on the tape.

If multiple disk volumes are dumped to a given TAPEX DD, ABR will create multiple files on the tape (or tape aggregate if more than one tape volume is used), one file for each disk volume, using the naming convention in [Section 50.01](#).

```
EXAMPLE: //TAPE1      DD   DSN=ABR1,UNIT=3590-1,DISP=(NEW,KEEP),
          //          VOL=( , , 255 ),LABEL=EXPDT=99000
```

DCB parameters are not required and should be omitted. However, tape unit hardware compaction (sometimes called IDRC, available on most tape cartridge drives) can be requested by adding DCB=TRTCH=COMP to your DD statement; Tape hardware compaction may be the default depending on local MVS options. For tapes attached by ESCON or FICON channels, Innovation recommends use of tape hardware compaction instead of FDR software compression (the COMPRESS= option).

50.04 CONTINUED . . .

If you are running the StorageTek ExHPDM (High Performance Data Mover) software product, you can direct ABR volume backups to ExHPDM with the SUBSYS=JCL operand, e.g.,

```
EXAMPLE: //TAPE1      DD  SUBSYS=(SOV,'CLASS(FDRBKUP)')
```

Please read [Section 80.33](#) and the ExHPDM program documentation for more details.

The IBM Virtual Tape Server (VTS) and similar products from other vendors are supported. In a VTS, data written to "tape" is really written to disk internal to the VTS and is later moved to high-capacity tapes such as Magstars, resulting in much better physical tape utilization. When a tape is required for input, the data is staged back to the internal disk.

LAST TAPE Option:

The LAST TAPE option of ABR allows you to add backup files to a tape or tape aggregate created by a previous ABR step (even if that step is in another job and even if it was run on a previous day). This option is controlled entirely through JCL. To request LAST TAPE, the TAPEX DD is similar to that described above except that you specify:

DSN=FDRABR.LASTAPE.xxxxxxxx – the last index is 1 to 8 characters of your choice; you may have multiple LASTAPE files for various purposes. This name will be cataloged in the ABR catalog to record the tape volume serial and file number where ABR is to start its output on the next run.

DISP=(MOD,KEEP) – this tells ABR to locate the FDRABR.LASTAPE.xxxxxxxx file in the catalog, verify that the file exists on the output tape, and begin outputting to the tape at that point. If the name is not cataloged, ABR will call for a scratch tape and begin at file 1, so you can force ABR to start on a new tape simply by uncataloging the LASTAPE name. Also, if you specify NEW instead of MOD, ABR will ignore the LASTAPE file and use scratch tapes (but it will still record the new LASTAPE file for future use).

VOL= – volume serials should not be specified, but the volume count should be given, e.g.,
VOL=(, , , 255).

```
EXAMPLE: //TAPE1      DD  DSN=FDRABR.LASTAPE.ABR1,
//                      UNIT=3490,DISP=(MOD,KEEP),
//                      VOL=( , , , 255 ),LABEL=EXPDT=99000
```

On cartridge tape drives, ABR will record the location of the LASTAPE file and will use hardware high-speed positioning when positioning to add files to the backup tape.

Disk Output:

ABR Volume Backups are not normally directed to disk files, but disk output is supported. See [Section 51.04](#) for details on specifying disk volumes for ABR output files. Since backups on disk cannot be restored by SAR (Stand-Alone Restore) and are of no use for disaster recovery, you may want to copy or move them to tape using the FDRTCOPY or FDRTSEL utilities documented in [Section 60](#).

WARNING: Tapes created by ABR cannot be copied using normal copy programs. Use the INNOVATION provided program, FDRTCOPY, to copy ABR tapes.

50.04 CONTINUED . . .

SNAP/SPLIT/PSPLIT:

If you are using FDR InstantBackup for your ABR volume backups, the ABR job will consist of two steps: the first will execute a SNAP, SPLIT or PSPLIT statement, and the second will execute the DUMP statement with a SNAP=, BCV= or PPRC= operand. Both the SNAP/SPLIT/PSPLIT step and the DUMP step must have the same TAPE_x and TAPE_{xx} DD statements, but the DDs in the SNAP/SPLIT/PSPLIT step will be DUMMY, e.g.,

```
//TAPE1 DD DUMMY
//TAPE11 DD DUMMY
//TAPE2 DD DUMMY
```

However, if you specify the EXPDT= or RETPD= operands on the TAPE_x/xx DD statements in the actual DUMP step, you must specify the same operands on the TAPE_x/xx DUMMY statements in the SNAP/SPLIT/PSPLIT step, e.g.,

```
//TAPE1 DD DUMMY,RETPD=250
```

Simulation:

If simulation of ABR Volume backups is requested by the SIM statement, you must specify:

```
//TAPE1 DD DUMMY
```

**TAPE_{xx} DD
STATEMENT**

Specifies that ABR is to create a duplicate backup (COPY 2). xx must specify the same character twice with xx corresponding to the x of the TAPE_x statement.

EXAMPLE: //TAPE11 DD DSN=FDR11,DISP=(,KEEP),UNIT=TAPE

will produce a tape which contains the same backup as

```
//TAPE1 DD DSN=FDR1,DISP=(,KEEP),UNIT=TAPE
```

NOTE: TAPE_{xx} may specify any of the options as documented for TAPE_x, including LASTAPE. However, a unique LASTAPE name must be used for each TAPE DD in the job step.

ABR will read the disk once and write the same data to TAPE_x and TAPE_{xx} concurrently. These are known to ABR as COPY 1 (TAPE_x) and COPY 2 (TAPE_{xx}).

**ABRBKDQ DD
STATEMENT**

Specifies the remote queue data set for BACKUP operations. This DD statement is optional. If specified, ABR will read the control statements contained within, if any, and append these statements to the SYSIN data set. The SYSIN data set must contain at least a DUMP TYPE=ABR, DSF or AUTO statement. After reading the control statements, ABR will reset the file to null (empty) data set except on SIM. DISP=SHR should be specified, since ABR internally serializes this data set.

CAUTION: This ABR execution must process all of the volumes which contain the data sets specified in the remote queue. This can easily be done by use of the ONLVOL operand on the DUMP statement.

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for ABR. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

50.05 VOLUME BACKUP DUMP/SIM/SNAP/SPLIT/PSPLIT STATEMENTS

DUMP	TYPE=FDR ABR DSF AUTO	,MAXAUTO=nnn
D	,AUTOUPD= <u>YES</u> NO	,MAXCARDS=nnnn
SIM	,BCV=(USE,RET)	,MAXDD=nnnn
SNAP	,BUFNO= <u>MAX</u> nn	,MAXERR=nnnn
SPLIT	,COMPRESS=ALL COPY1 COPY2	,MAXFILE=nnnn
PSPLIT	,COPY1=COPY2	,NOINIT
	,DATA=ALL USED	,ONLINE
	,DATEP=NONE	,ONLVOL
	,DCT=YES <u>NO</u>	,PPRC=(USE,RET)
	,DD=ALL	,PRINT=DSN ABR
	,DSNENQ= <u>NONE</u> TEST USE HAVE	,RETPD=dddd
	,ENQ=OFF <u>ON</u> RESERVE	,RETPD2=dddd
	,ENQERR=NO	,RTC=YES <u>NO</u>
	,ENQERR=BYPASS <u>PROCESS</u>	,SELTERR=NO <u>YES</u>
	,FORMAT= <u>NEW</u> SPLIT	,SMSCONSTRUCT= <u>YES</u> NO
	,HFS=QUIESCE	,SMSMANAGE= <u>NO</u> YES
	,ICFCORE=nnnnnn	,SNAP=(USE,REL)
		,UPDATEFLAG=NOCHANGE
		,VOLSORT= <u>YES</u> NO

DUMP STATEMENT The DUMP statement initiates a Volume Backup of disk packs or data sets. Only one DUMP statement is allowed per execution of ABR.

If TYPE=FDR is specified, a full-volume backup is taken of each selected disk volume.

If TYPE=ABR, AUTO or DSF is specified, an incremental backup is taken of each selected disk volume. In some circumstances, a full-volume backup may be taken in place of the incremental. [See Section 50.01](#) for details.

SIM STATEMENT If SIM is specified instead of DUMP, ABR will perform the DUMP function in a simulation mode. The TAPEX DD Statement must specify DD DUMMY. ABR will not perform the actual DUMP operation, but it will print the data set names selected by this operation using the PRINT VTOC format. Only one SIM statement is allowed per execution of ABR.

SNAP STATEMENT The SNAP statement is used with FDR InstantBackup when the disk subsystem is an IBM RVA (Ramac Virtual Array) or StorageTek Iceberg or SVA (Share Virtual Array), with the optional Snapshot feature. [Section 26](#) contains details of its use. However, the volume backup will consist of two steps, a SNAP step and a DUMP step. You should specify all of the same operands on the SNAP and DUMP statements, except for the SNAP= operand which is added to the DUMP step to

50.05 CONTINUED . . .

invoke FDR InstantBackup.

SPLIT STATEMENT The SPLIT statement is used with FDR InstantBackup when the disk subsystem is an EMC Symmetrix with the optional Timefinder feature. [Section 25](#) contains details of its use. However, the volume backup will consist of two steps, a SPLIT step and a DUMP step. You should specify all of the same operands on the SPLIT and DUMP statements, except for the BCV= operand which is added to the DUMP step to invoke FDR InstantBackup.

PSPLIT STATEMENT The PSPLIT statement is used with FDR InstantBackup when the disk subsystem is an HDS 7700E with the optional ShadowImage feature. [Section 27](#) contains details of its use. However, the volume backup will consist of two steps, a PSPLIT step and a DUMP step. You should specify all of the same operands on the PSPLIT and DUMP statements, except the PPRC=operand which is added to the DUMP step to invoke FDR InstantBackup.

OPERANDS TYPE= Specifies the type of dump to be performed. It is required.\

FDR – Specifies that a full volume backup of all the volumes selected will be taken. This begins a new generation of backups for those disk volumes.

ABR – Specifies that an incremental backup of all the volumes selected will be taken. ABR examines the update flag set by the operating system. Any data set which has been updated or created since the last Volume Backup will be dumped, as well as any data set for which backup has been requested through the backup remote queue or by a SELECT statement. This function creates a new cycle within the current generation. [See Section 50.02](#) for more details.

DSF – Specifies that a manual backup of only the data sets specified by SELECT statements or the DD=ALL option will be taken. Note that if an ABR full-volume restore is performed ([see Section 50.10](#)) and the most recent incremental backup read is a TYPE=DSF backup, it will be bypassed and the data sets on the backup will not be restored from that backup.

AUTO – Specifies that ABR is to automatically determine whether a full volume or incremental backup is to be taken. ABR compares the number of incremental backups taken with AUTOUPD=YES ("auto cycles") in effect with the maximum number of cycles specified in the ABR model DSCB on the volume. When the number of auto cycles exceeds the cycle limit, a TYPE=FDR full volume dump is taken.

The default maximum cycles before a full volume backup is forced is 10 but can be changed by program FDRABRM or the ABR ISPF Dialogs.

AUTOUPD=

NO – specifies that ABR will not increment the auto cycle count for this execution. This is useful if this is a special run or a run outside of auto control. May be specified with TYPE=ABR, DSF or AUTO.

YES – specifies that this incremental backup will count for TYPE=AUTO operations.

NOTE: ABR maintains two cycle numbers. The CYCLE (last two digits of the backup name) is incremented on each backup run. The AUTO CYCLE is incremented if AUTOUPD=YES is specified or defaulted.

Default is YES.

50.05 CONTINUED . . .

BCV=

This operand is used only if you are licensed for FDR InstantBackup and only if the disk being backed up is in an EMC Symmetrix disk subsystem with the Timefinder option. It is specified on a DUMP statement following a step which executed an ABR SPLIT statement to split a BCV from its online volume, creating a frozen point-in-time volume image.. Possible values are:

USE - move the point-in-time backup from the BCV to tape.

(USE,RET) – same as BCV=USE but at the end of moving each BCV to tape, ABR re-establishes the BCV so that it is again mirroring the online volume

Please read [Section 25](#) for complete documentation on FDR InstantBackup and Timefinder.

BUFNO=

specifies how many buffers will be used for dumping each disk volume. Each buffer holds one disk track. The buffers acquired will be divided into 2 sets in order to overlap input and output I/O operations; each disk I/O will read disk tracks into one half of the buffers.

MAX – buffers sufficient to read 1 cylinder of the input disk are acquired.

nn - the specified number of buffers is acquired.

Default: MAX. **Innovation recommends that you do not override the default.** However, BUFNO=2 will be forced when a backup (output) data set is on disk.

COMPRESS=

Controls the use of FDR software compression. Values for COMPRESS= are:

ALL – the backup file for both copies (TAPEX and TAPEXX) is to be compressed.

COPY1 – only the backup on TAPEX DD statements will be compressed.

COPY2 – only the backup on TAPEXX DD statements will be compressed.

See "Memory Requirements" in [Section 50.03](#) for the additional storage required by COMPRESS=. COMPRESS= is ignored if DCT=YES is also specified.

Default: backups will not be compressed.

COMPRESS is recommended for backups to disk files, and for tape backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of IDRC (tape hardware compression) usually results in better performance.

COPY1=COPY2

Swaps the copy numbers created on the TAPEX and TAPEXX DD statements; COPY 2 will be created on the TAPEX DD and COPY 1 on the TAPEXX DD (see [Section 50.01](#) for the backup data set naming conventions including the copy number). If only TAPEX is present, only COPY 2 will be created; if you send COPY 2 offsite, this can be useful when you wish to create *only* an offsite backup (no COPY 1 for onsite recovery).

50.05 CONTINUED . . .

DATA=	<p>USED - only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be backed up. On most volumes, this will make the dump run faster.</p> <p>ALL - all allocated tracks of all selected data sets will be backed up. You may need to specify DATA=ALL if the data sets to be backed up include JES2 spool data sets or CICS log data sets, since they usually do not have valid last block pointers.</p> <p>Default: ALL for full-volume backups, USED for incremental (data set) backups.</p>
DATEP=	<p>NONE -- specifies that ABR will not check the expiration date of the most recent full-volume dump.</p> <p>Default is that ABR will force a full volume backup if the most recent full volume backup of the disk being processed has reached its expiration date, as recorded in the ABR Model DSCB on the volume.</p>
DCT=	<p>DCT= is valid only if you are licensed for FDR InstantBackup. It will be honored only if the disk being backed up is in a disk subsystem with the HSDM option (High Speed Data Mover). HSDM allows FDR to backup and restore the internal compressed images of disk tracks, improving backup elapsed times up to 60%. It can also be specified as DUMPCOMPRESSEDTRACK=.</p> <p>YES -- use HSDM for any volume where the disk hardware has the HSDM feature installed. Normal backup will be used for other volumes.</p> <p>NO -- do not use HSDM.</p> <p>Default is NO.</p> <p>Note that DCT=YES implies RTC=YES; see the description of RTC= for its benefits.</p>
DD=ALL	<p>Specifies that ABR is to scan all the DISKxxxx DD Statements for the DSNAMES coded. ABR will attempt to backup these data sets. ABR will combine all of the data sets from one volume into one backup file. This feature is useful if data sets are allocated with non-specific volume requests or by SMS or other allocation software.</p>

50.05 CONTINUED . . .

DSNENQ=

Specifies whether all of the data sets selected for backup will be ENQed. See “Data Set Enqueue Option” in [Section 50.03](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a warning message is issued for the data set but the data set will still be dumped unless the ENQERR=BYPASS operand is specified. A successful ENQ will prevent any other task from using the data set until the backup of that volume is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be dumped. The options for DSNENQ= are:

USE – The data sets will be enqueued for the duration of the backup from this disk volume. This is the most frequently used option.

TEST – The data sets will only be tested to see if they are enqueued to another task at the time that the dump from this volume starts.

HAVE – The data sets will be enqueued for the duration of the dump. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not ENQ the data set)

RETRY (try the ENQ again)

NONE – No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: NONE.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the backup is complete. You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements..

ENQ=

Specifies whether an ENQ should be done on the VTOC of each disk volume while data sets from it are being backed up. See “VTOC Enqueue Option” in [Section 50.03](#) for more details.

ON – the VTOC of each disk volume will be ENQed during its backup. This ENQ may be effective only on the system where the backup is executing; other systems may still be able to update the VTOC.

RESERVE – in addition to the ENQ, a hardware RESERVE will be issued on each disk volume during its backup. This is meaningful only on a system with “shared DASD” where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF – the VTOC will not be enqueued or reserved during the backup.

Default is ON.

50.05 CONTINUED . . .

- ENQERR=** **NO** – If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see “Step Termination” in [Section 50.03](#)). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally. Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.
- ENQERR=** Specifies processing if the DSNENQ= option finds that a data set is in use (enqueued):
- BYPASS** – do not backup an active data set. This is ignored for a full-volume backup.
- PROCESS** – backup a data set even if it is active (a warning message will still be produced). Default: PROCESS.
- NOTE:** both ENQERR=NO and ENQERR=BYPASS/PROCESS may be specified on the same DUMP statement.
- FORMAT=** Specifies the format of the sequential backup file.
- NEW** – the backup will be created using a maximum of a 56K blocksize. A block will contain the image of one or more tracks from the input disk.
- SPLIT** – the backup will be created using a maximum blocksize of 32K. For blocks that would be more than 32K, they are written as 2 blocks of 32K or less. However, FORMAT=SPLIT causes use of a disk I/O technique which is less efficient than that used by FORMAT=NEW, which will impact backup performance.
- WARNING: If you use a normal copy program (ex: IEBGENER) to copy a backup file created with FORMAT=NEW, you will not get any error messages, but the resulting tape will not be usable for a restore. Tapes in the new format must only be copied with the utilities FDRTCOPY and FDRtsel (See Section 60).***
- Default: NEW if backup on tape – SPLIT if backup is on disk.
- HFS=** **QUIESCE** invokes special processing when HFS (Hierarchical File System) data sets are backed up. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to Unix System Services, so a "quiesce" call is issued to prevent updates to the HFS data set during the backup. If the quiesce fails and ENQERR=BYPASS was specified, the HFS data set will not be backed up by incremental backup. See "HFS" in [Section 80.11](#) for details and security requirements
- Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for **all** data sets being backed up, not just HFS data sets.
- Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table ([see Section 90](#)). If you use the default you should unmount the HFS file system before the backup to be sure of getting a usable backup.

50.05 CONTINUED . . .

- ICFCORE=** Specifies that the size of the table used to store the ICF VSAM cluster and component names. The backup must save all of the component names and their associated clusters which exist on the current input disk volume, in order to match up VTOC DSCBs (with the component name) to cluster names for selection. nnnnnn is specified in bytes and must be large enough to contain all the VSAM names (depending on the length and number of names).
- NOTE:** Specifying ICFCORE= will increase the backup region requirement by the value specified. The default value uses one of the dump buffers and imposes no additional memory requirement.
- Default: 53248, which will hold about 650 components.
- MAXAUTO=** Used with DUMP TYPE=AUTO only. Specifies the maximum number volumes for which full volume dumps can be forced by TYPE=AUTO within this execution. Incremental backups will continue to be taken for any additional volumes, even if they have exceeded their auto cycle limit, unless full-volumes are forced by the last full-volume expiration date or maximum cycle number (63).
- Default is that a full-volume backup will be forced for all volumes exceeding the auto cycle limit.
- MAXCARDS=** Enables ABR to accept additional SELECT, EXCLUDE and MOUNT statements during this execution. You can specify values up to 65535, but values over 100 will increase the region required by ABR.
- Default is 100 SELECT, EXCLUDE and MOUNT statements.
- MAXDD=** Specifies the maximum number of disk volume serials that ABR can process in this execution. ABR builds its volume list from DISKxxxx DD statements and/or ONLINE volumes if ONLINE, ONLVOL or MOUNT is coded. Once this limit is reached additional volumes will be ignored.
- Default is 256 volumes, unless overridden by the MAXONLINE option in the FDR Global Option Table ([see Section 90](#)).
- MAXERR=** The number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup was written to the ExHPDM subsystem, MAXERR=1 is recommended.
- WARNING:** MAXERR over the default value may result in the loss of many data sets. This option should only be used when necessary and with care.
- Default: 20 errors.

50.05 CONTINUED . . .

MAXFILE=	<p>Specifies the maximum number of files ABR will create on tape. May specify from 1 to 4095. When the maximum file number is exceeded, ABR will start a new tape using file sequence number 1. A MAXFILE= over 255 may be appropriate when you are outputting to high-capacity tapes such as IBM 3590 Magstar or StorageTek Redwood or 9840.</p> <p>Default is 255 unless overridden in the FDR Global Option table (See Section 90).</p>
NOINIT	<p>Valid only on a SIM statement. Specifies that ABR is to perform the simulation on the volumes specified even if they are not initialized for ABR processing. ABR error messages will be issued but the simulation will be done.</p>
ONLINE ONLVOL	<p>If these operands are omitted, ABR will only process volumes specified by DISKxxxx DD statements or ABR MOUNT statements.</p> <p>ONLINE – specifies that ABR is to process every volume that is ONLINE to the system; if there are more than 255 online volumes, you will need to specify MAXDD= to increase that limit. With ONLINE, there is usually no need to specify DISKxxxx DD statements or MOUNT statements. However, if you have certain disk volumes which you want processed first, specify them on DISKxxxx DD statements; other online volumes will be appended to the list created from the DD statements.</p> <p>ONLVOL – specifies that ABR is to scan all SELECT statements for the operands VOL=, VOLG=, or CATDSN= (for volume selected from the catalog). These volumes, if online, will be appended to the ABR volume list.</p> <p>If ABR finds that the user specified EXCLUDE statements with ALLDSN and VOL or VOLG, these volumes will not be included in the ONLINE list unless a DISKxxxx DD statement is present for this volume. This is a way of excluding certain volumes, for example, EXCLUDE ALLDSN, VOLG=WORK.</p>
PPRC	<p>This operand is used only if you are licensed for FDR InstantBackup and only if the disk being backed up is HDS 7700E disk subsystem with the ShadowImage option. It is specified on a DUMP statement following a step which executed an ABR PSPLIT statement to split a ShadowImage copy from its online volume, creating a frozen point-in-time volume image. Possible values are:</p> <p>USE - move the point-in-time backup from the ShadowImage copy to tape.</p> <p>(USE,RET) – same as PPRC=USE but at the end of moving each ShadowImage copy to tape, ABR re-establishes the ShadowImage copy so that it is again mirroring the online volume</p> <p>Please read Section 27 for complete documentation on FDR InstantBackup and ShadowImage.</p>

50.05 CONTINUED . . .

- PRINT=** Requests that ABR print the names of all data sets selected for backup.
- DSN** – Specifies that ABR is to simply print the data set names.
- ABR** – Specifies that ABR is to use the PRINT VTOC format of program FDRABRP ([See Section 53](#)) for each data set that was dumped.
- Default: DSN except for full-volume backups where data set names are not printed by default. PRINT=ABR is the default for a SIM statement (simulation).
- RETPD=** **RETPD=** specifies the number of days (1 to 9999) that COPY1 (TAPEx) backups will be kept. ABR will calculate an expiration date from this value. If RETPD2= is omitted, this same expiration will apply to all COPY2 backups created in the same step.
- RETPD2=** specifies the number of days (1 to 9999) that COPY2 (TAPExx) backups will be kept. ABR will calculate an expiration date from this value.
- [See Section 50.01](#) for a discussion of tape retention.
- Default for a full-volume backup (TYPE=FDR) is the default stored in the ABR Model DSCB on this volume, set by the ABR utility FDRABRM or the FDR Volume Initialization panel (A.I.8); separate defaults for RETPD= and RETPD2= are stored in the model. For incremental backups, the default is the expiration date of the most recent COPY1 or COPY2 full-volume backup as recorded in the ABR Model DSCB. These defaults can be overridden by RETPD= or EXPDT= on individual TAPE DD statements
- RTC=** It can also be specified as READTRACKCCW=.
- YES** – use READ TRACK CCWs to read disk data tracks. RTC=YES also causes:
- up to 1 cylinder of disk data is read at a time.
 - FDR buffers are moved above the 16MB line (about 2MB per concurrent backup), allowing more concurrent backups to be run in one step.
 - the elapsed time of ABR backups when the backup data set is itself on disk is significantly improved.
- NO** – use other CCWs to read disk data tracks.
- Default is NO.

50.05 CONTINUED . . .

SELTERR=

Specifies what will happen at step termination if ABR finds that a SELECT or EXCLUDE statement was never referenced (no data set on any input disk was selected by the statement):

NO – a condition code or ABEND is **not** to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.

YES – a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.

Default: YES unless overridden in the FDR Global Option Table (See [Section 90](#)).

SMSCONSTRUCT=

YES – ABR is not to process SMS-managed volumes unless their associated storage group has the attribute AUTO-DUMP=YES and AUTO-BACKUP=YES. (Volumes with AUTO-DUMP=YES but AUTO-BACKUP=NO will be selected for TYPE=FDR runs only, and bypassed on all incrementals). The SMS volumes must still have a proper ABR Model DSCB.

NO – ABR will bypass selected volumes based only on the presence or absence of the ABR Model DSCB, ignoring the SMS storage class attribute.

Default is YES.

SMSMANAGE=

NO – ABR is not to exclude data sets from SMS-managed volumes based on the attributes of their SMS management class.

YES – data sets on SMS-managed volumes will be excluded from data set backups if their associated SMS management class has the attribute "ADMIN OR USER COMMAND BACKUP=NONE" (for all data set backups) or "AUTO-BACKUP=NO" (for TYPE=ABR/AUTO). If not excluded by the above, data sets will be selected by normal ABR rules (See [Section 50.02](#)) or by SELECT statements. See [Section 70](#) for more details on SMS processing.

Default is NO.

WARNING: if data sets are excluded from incremental backups because of AUTO-BACKUP=NO, an ABR full-volume recovery will not properly recover the latest copy of those data sets.

50.05 CONTINUED . . .

SNAP=

This operand is used only if you are licensed for FDR InstantBackup and only if the disk being backed up is in an IBM RVA or StorageTek Iceberg or SVA disk subsystem with the Snapshot option. It is specified on a DUMP statement following a step which executed an ABR SNAP statement to snap an online volume to an offline target volume, creating a frozen point-in-time volume image. Possible values are:

USE - move the point-in-time backup from the snapshot target volume to tape.

(USE,REL) - same as SNAP=USE, but at the end of backing up each snapped volume, ABR will issue a request to delete all the storage assigned to the snapped copy (Deleted Space Release), except for the label track (cylinder 0 track 0). This is recommended since it keeps the NCL (Net Capacity Load) in the disk subsystem down by releasing the tracks of the snapped copy as soon as they are no longer needed.

Please read [Section 26](#) for complete documentation on FDR InstantBackup and Snapshot.

UPDATEFLAG=

NOCHANGE – ABR will not reset the update/change flag in the Format 1 DSCBs of data sets it backs up in this execution. Since the update flag remains set, the data sets will be included in the next backup job even if they are not updated again.

UPDATEFLAG=NOCHANGE is useful when you are testing the operation of ABR compared to another DASD management product, such as IBM's HSM. If you execute ABR Volume Backups first, followed by HSM backups, ABR will backup all updated data sets and leave the update flag on. Then HSM will backup the same data sets and turn the update flag off. This allows each ABR and HSM backup to select the same data sets and do equivalent work so that you can compare elapsed time and resources used by both.

VOLSORT=

YES — if you have more than one TAPEX DD statement in this ABR step, ABR will sort the volumes to be backed up by the last digit of their MVS device address, to attempt to balance channel and control unit utilization during concurrent backups. In other words, all volumes whose MVS address is xxx0 will be processed first, then xxx1, etc. This only applies to volumes selected by the ONLINE or ONLVOL operands or MOUNT statements.

NO — volumes will be processed in the order that they are found during a scan of the system UCBs, which is usually (but not always) in MVS device address order.

Default is YES except that NO is forced if you have only one TAPEX DD statement.

The above does not apply to volumes selected by DISKxxxx DD statements which will always be processed first, in the order the DD statements appear in the JCL.

50.06 VOLUME BACKUP SELECT, EXCLUDE AND MOUNT STATEMENTS

SELECT	DSN=filter	,DSNENQ=NONE
S	CATDSN=dsname	
	DD=ddname	,DSORG=(xx,xx..)
EXCLUDE	ALLDSN	
X		,PRTALIAS
	,CATALOG=catname	
	,MCATALOG=catname	,UPDATE
	,CATLIMITGDG=n	,VOL=vvvvvv
		,VOLG=vvvvvv
	,DATA=ALL NONE	
MOUNT	VOL=vvvvvv VOLG=vvvvvv STORGRP=storagegroup	
	,BCVUNIT=uuuuu	
	,SNAPUNIT=uuuuu	

**SELECT/
EXCLUDE
STATEMENTS**

These statements select additional data sets to be backed up. They are ignored for full-volume backups except for EXCLUDE ALLDSN statements used to exclude whole volumes. For incremental backups (DUMP TYPE=ABR or AUTO) data sets to be backed up are selected by the rules in [Section 50.02](#) and SELECT statements are normally not required. However you may use SELECT statements to include data sets which would not normally be included in an incremental backup. For DUMP TYPE=DSF, SELECT statements are required to identify the data sets to be included.

A **SELECT** statement identifies an individual data set or group of data sets to be processed. ABR will look for these data sets on every volume processed unless the VOL= or VOLG= operand is included; if CATDSN= is used, only the volume identified from the MVS catalogs will be checked for the data sets.

An **EXCLUDE** statement identifies data sets from within those selected by SELECT statements which are not to be processed; it can also be used to exclude data sets which would otherwise be included in an incremental backup by the rules in [Section 50.02](#).

All data sets in the VTOCs of DASD volumes processed by ABR will be compared to these statements to identify those to be processed; each data set will be compared to each control statement until a match is found. It must match **all** criteria specified on the statement to qualify, e.g., if DSN= and DSORG= are both specified, it must be a data set with the right name, having the indicated organization. A maximum of 100 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since ABR will only dump data sets which are selected, EXCLUDE statements can be used to exclude certain data sets from within a larger group on a SELECT statement

50.06 CONTINUED . . .

**SELECT/
EXCLUDE
STATEMENTS
(CONTINUED)**

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A.B.**
SELECT  DSN=A.**
```

Example 2. Select all data sets except partitioned (PDSs):

```
EXCLUDE ALLDSN,DSORG=PO
SELECT  ALLDSN
```

You can also use an EXCLUDE ALLDSN statement with VOL= or VOLG= to specify entire volumes which are not to be processed by ABR. This works even with full-volume backups. This is a way to exclude certain volumes that would otherwise be selected by the ONLINE or ONLVOL operands or MOUNT statements.

**MOUNT
STATEMENTS**

The **MOUNT** statement is used to specify additional volumes to be processed by ABR. VOL=, VOLG=, or STORGRP= must be specified. STORGRP= can be used only on systems with SMS active; it will select all volumes in the specified SMS storage group. MOUNT statements should follow SELECT/EXCLUDE statements, if any.

**DUMPING ICF
VSAM FILES**

ICF VSAM clusters can be selected by specifying the fully-qualified base cluster name or matching on the base cluster with generic data set selection. When selected, all components of that cluster that exist on the volumes being processed will be dumped or printed, including alternate indexes and key range components. ABR will not examine ICF VSAM component names when processing SELECT/EXCLUDE statements; components will be selected only if their cluster name is selected. For further information, [see Section 80.13](#) "VSAM Special Considerations".

OPERANDS DSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the VTOCs of selected volumes.

```
Examples : DSN=USER1.JCL.CNTL
           DSN=**LIST
           DSN=PROD++.**.LIB*
```

DSN= cannot select GDGs by relative generation number; use DD= for that purpose.

50.06 CONTINUED . . .

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= operand (you should specify the ONLVOL operand so that those volumes are automatically added to the ABR volume list for processing). Specification of a relative generation number for GDG data sets is supported, e.g., CATDSN=A.B(-1)

If a filter is specified, then catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT CATDSN=dsname was present for each of them. It may be necessary to specify MAXCARDS=nnnnn if a large number of data sets are selected by the filter.

Additional considerations for CATDSN=filter are explained in [Section 80.14](#).

CATDSN= is supported only on SELECT statements. However, a preceding EXCLUDE statement with DSN= and/or VOL= can exclude data sets from selection by CATDSN=.

If the VOL= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

Examples: CATDSN=USER1.JCL.CNTL
 CATDSN=**MASTER(0)
 CATDSN=PROD++.**.LIB*

Normally CATDSN= will not display the data sets it selects from the catalogs, you will see the names only when ABR actually finds and selects the data sets in the VTOCs of the volumes they are cataloged to. To display all of the data sets selected specify PCATDSN=filter.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

DD=

Specifies that a data set name is to be taken from a DD statement. The value is a the DDNAME of a JCL statement in the ABR step. Using this option enables you to specify a non-standard data set name or a generation data set (GDG) relative generation.

Example: SELECT DD=DD1
 //DD1 DD DSN=A.B.C(0),DISP=SHR

Note that although DD= copies the data set name from the DD statement, it does not use the volume pointed to by that DD unless the DDNAME is DISKxxxx. The data set will be searched for only on the volumes processed by ABR; the volume containing the data set must be among them.

50.06 CONTINUED . . .

ALLDSN	<p>Specifies that ABR is to backup all data sets on the volumes specified. If TYPE=DSF is specified, SELECT ALLDSN will not dump the VTOC and label track.</p> <p>EXCLUDE ALLDSN with the VOL= or VOLG= operand will totally exclude those volumes from ABR processing.</p> <p>NOTE: DSN=, CATDSN=, DD= and ALLDSN are mutually exclusive. One and only one DSN=, CATDSN=, DD= or ALLDSN operand must be specified on a SELECT or EXCLUDE Statement.</p>
BCVUNIT=	<p>Can only be used on a MOUNT statement following a SPLIT statement and must be preceded by a VOL= operand, e.g.,</p> <pre>MOUNT VOL=PROD01,BCVUNIT=07CA</pre> <p>VOL= specifies an online volume and BCVUNIT= specifies the 4-digit MVS device address of the BCV currently assigned to the online volume. BCVUNIT= is optional; if not specified, FDR will locate the BCV address from information provided by the EMC Symmetrix. However, if your Symmetrix has multiple non-contiguous MVS address ranges assigned, this may require a lengthy search through many offline disk devices.</p>
CATALOG= MCATALOG=	<p>Specifies the name of a user catalog (CATALOG=) or alternate master catalog (MCATALOG=) to search when CATDSN= is specified. See Section 80 for details.</p> <p>Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter.</p>
CATLIMITGDG=	<p>May be used with CATDSN=filter to limit the selection of GDGs from the catalogs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:</p> <p>n will cause only the most recently created "n" generations to be selected.</p> <p>-n will cause only generation (-n) to be selected.</p> <p>Default is that all the generations of selected GDGs will be selected unless a relative generation number is specified at the end of the filter, e.g., CATDSN=filter(-2).</p>
DATA=	<p>ALL – specifies that ABR will backup the entire allocated space of the selected data sets. Normally it will only process up to the last block pointer (end-of-file) on input PS or PO data sets for incremental (data set) backups. Should be used if the last block pointer of certain data sets is invalid.</p> <p>NONE – no data tracks will be backed up for any type of data sets selected. The backup data set will consist only of the FDR control records necessary to allocate and update the characteristics of the output data sets during a restore. A DATA=NONE backup can be used to allocate all the data sets on the backup on output volumes, as long as the data in those data sets will be recovered by other means (such as data base backups). Member DATANONE in the FDR ICL (Installation Control Library) has more information on the use of DATA=NONE.</p>

50.06 CONTINUED . . .

DSORG=	<p>Specifies that these data sets is not to be selected unless DSORG matches one of the DSORGS specified. If more than one DSORG is specified, they must be enclosed in parenthesis.</p> <p>VALID DSORGS are:</p> <p>DA – BDAM PS – SEQUENTIAL AM – ALL VSAM EF – ICF VSAM IS – ISAM PO – PARTITIONED UN – UNDEFINED UM – UNMOVABLE</p>
DSNENQ=	<p>NONE – the data set enqueue will not be done for the selected data sets. This can be used to override the DSNENQ= operand on the DUMP statement for certain data sets that you know will probably be ENQed by another task.</p> <p>Default: determined by the DSNENQ option on the DUMP statement.</p>
PRTALIAS	<p>When used on a SELECT statement with CATDSN= will display all of the alias names and user catalogs that were searched. This is effective only when the data set name mask you provided forces CATDSN= to start in the master catalog and search one or more alias catalogs. If the non-mask characters at the beginning of the mask are sufficient to cause CATDSN= to begin its search in a user catalog, PRTALIAS is ignored.</p>
SNAPUNIT=	<p>Can only be used on a MOUNT statement following a SNAP statement and must be preceded by a VOL= operand, e.g.,</p> <pre>MOUNT VOL=PROD01, SNAPUNIT=07CA</pre> <p>VOL= specifies the online volume to be snapped and SNAPUNIT= specifies the 4-digit MVS device address of the target volume to which it is to be snapped.</p>
STORGRP=	<p>Can only be used on the MOUNT statement, and only on systems with SMS (System Managed Storage) active. Will select all online volumes in the specified SMS storage group.</p> <p>Example: MOUNT STORGRP=DBLARGE</p> <p>will cause ABR to process all volumes in that storage group.</p>
UPDATE	<p>Specifies that the data sets or ICF VSAM clusters that match this control statement will only be dumped if the update indicator is set in their Format 1 DSCB, indicating that the data set has been updated since the last ABR Volume Backup.</p> <p>NOTE: If you wish to do incremental backups of only the specified data sets (using the update indicator) specify TYPE=DSF.</p> <pre>Example: DUMP TYPE=DSF, COMPRESS=ALL SELECT DSN=USER**, UPDATE SELECT DSN=XYZ**, VOLG=VVVVV, UPDATE</pre>
VOL=	<p>Specifies a volume serial number (up to 6 characters) to which this statement applies. If not specified, ABR will examine all volumes processed for the data sets named on this statement. For MOUNT statements, it will add the named volume to the ABR volume list if it is online.</p>

50.06 CONTINUED . . .

VOLG=

Specifies a volume serial prefix (1-5 characters) to which this statement applies. If not specified, ABR will scan all volumes processed for the data sets named on this statement. For MOUNT statements, it will process the all online volumes which start with the prefix.

NOTE: VOL=, VOLG= and STORGRP= are mutually exclusive. For MOUNT statements, one of those operands must be specified. If ONLVOL is specified on the DUMP statement, the volumes will be appended to the volume list.

50.07 RESTORE JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform restore from ABR Volume Backups:

**STEPLIB or
JOB LIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRABR), region requirement (REGION=), and optional PARM= operand. The minimum region required is 512K. However, some restore options, especially logical restore, may increase the region requirement, so a value of 1M or 2M is recommended; REGION=0M can be specified to get the largest below-the-line region available.

If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid RESTORE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=FDR,CONFMESS=NO'
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=ABR/ SELECT DSN=A.B.C'
```

If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**STEP CAT or
JOB CAT DD
STATEMENT**

For non-VSAM data sets, when FDR must catalog non-VSAM data sets, they will be cataloged in that user catalog instead of in the system catalog with the matching alias. This might be useful when you are creating a test system, to catalog test copies of production data sets in a test catalog.

However, there is one important exception: if the data set being cataloged is a GDG generation, the STEP CAT/JOB CAT must contain a GDG base for that GDG; if not, it will ignore the STEP CAT/JOB CAT and catalog into the regular aliased catalog, possibly deleting other valid generations.

For VSAM clusters, the target catalog is controlled by the ICFCAT= operand, described in [Section 50.08](#). With the proper ICFCAT= option, the STEP CAT/JOB CAT may be honored.

STEP CAT/JOB CAT should not be used if any data sets being restored are SMS-managed.

**SYS PRINT DD
STATEMENT**

Specifies the output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**FDR SUMM DD
STATEMENT**

(Optional) if present, ABR will write one-line messages for each volume restored, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set. It is used only for full-volume restores, not data set restores.

**SYS DUMP DD
STATEMENT**

Specifies the abend data set. Usually a SYSOUT data set. A SYS DUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

50.07 CONTINUED . . .

**TAPEx DD
STATEMENT**

Allocates an input tape drive to be used for the RESTORE operation. "x" may be 1 alphabetic (A-Z), numeric (0-9) or national (@ \$ in the US) characters. The national character # or its international equivalents is reserved for ABR dynamic allocation. This tape drive must be capable of reading all the tapes needed for this restore; if the tapes are in varying formats, such as 3480 and 3490E, the restore will fail (see "DYNTAPE note" below).

Although this DD allocates a tape drive, ABR will internally fill in the data set name, volume serials, and file number before it OPENS each tape file, so you do not have to provide any information about the backups to be read. However, MVS requires that you provide:

DSN= provide any dummy data set name. It will not be used by ABR, but MVS will ENQ on this name so multiple ABR restore jobs must use unique names.

VOL=SER= provide a dummy volume serial.

UNIT=(xxxx,,DEFER) this allocates the proper tape drive type (xxxx) but does not try to mount the dummy volume serial.

DISP=(OLD,KEEP) **required**

Only the first TAPEx DD provided is used by ABR during a RESTORE operation.

```
EXAMPLE: //TAPE1      DD    DSN=FDR,VOL=SER=FDR,
           //          UNIT=(TAPE,,DEFER),DISP=(OLD,KEEP)
```

If the backup to be restored was directed to the StorageTek ExHPDM software product as described earlier, you must specify the SUBSYS= operand on the TAPE DD statement to invoke ExHPDM and restore that backup, e.g.,

```
//TAPE1      DD    SUBSYS=SOV
```

However, when restoring with the DYNTAPE option, ABR will dynamically allocate ExHPDM backups with the SUBSYS=SOV parameter; contact Innovation if your ExHPDM subsystem name is not SOV.

NOTE: If the TAPEDD=x operand is used on a SELECT statement, the "x" must match a TAPEx DD statement, and ABR will read the backup data set specified on this DD to find the selected data sets. In this case, the TAPEx DD statement must identify an actual backup data set. If it is cataloged, at least the data set name and DISP=OLD are required; if not, the unit, volume serial and file sequence must be included.

DYNTAPE NOTE: If DYNTAPE is specified on the RESTORE statement, this DD statement is not used and can be omitted. ABR will dynamically allocate a TAPE# DD statement for the backup device. DYNTAPE should be used if the backup is on disk, in an automated tape library (ATL or silo) or if a mixture of tape device types must be read. If the tape volume is required for the full-volume recovery of two or more disk volumes in a row in the same ABR step, ABR will prevent the tape from unloading between restores. Do not use DYNTAPE and TAPEDD= in the same execution.

SIMULATION: If SIMREST is coded, this DD usually specifies DUMMY.

**DISKxxxx DD
STATEMENT**

For ABR full-volume restore (RESTORE TYPE=FDR) a DISKxxxx DD must be provided for each disk being restored unless the ONLINE operand is specified on the RESTORE statement

```
EXAMPLE: //DISK1234 DD    UNIT=SYSDA,VOL=SER=DSK123,DISP=OLD
```

For data set restores, DISKxxxx DDs are not used; ABR will dynamically allocate any required output volumes. Even if DISKxxxx DDs are present, they will not influence ABR's choice of output volumes (see [Section 50.08](#)).

50.07 CONTINUED . . .

ABRREST DD STATEMENT Specifies the remote queue data set for restores from the backup subsystem. This data set is optional. If specified, ABR will read the control statements contained within, if any, and append these statements to the SYSIN data set. The SYSIN data set must contain at least a RESTORE TYPE=ABR statement. If the operation is not a data set restore from the backup system, this file will be ignored. After reading the control statements, ABR will reset the file to null (empty) data set except on SIMREST.

NOTE: The TSO/ISPF panels ([Section 50.30](#)) or program FDRABRUT ([Section 51.40](#)) write to these remote queue data sets. DISP=SHR should be specified for ABRREST, since ABR internally controls access to this data set.

SYSIN DD STATEMENT Specifies a data set containing the control statements for ABR. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

50.08 DATA SET RESTORE STATEMENT

RESTORE R	TYPE=ABR	,ICFCAT=ORIGINAL STEP CAT ALIAS
	,BLKF=nn	,MAXCARDS=nn
SIMREST	,BYPASSACS	,NOCAT
	,BYPASSSMS	,RECAT
	,CATIFALLOC	,OPERATOR
	,COPY=n	,PRESTAGE
	,DATA=ALL	,RLSE
	,DD=ALL	,%FREE=nn
	,DSNENQ=NONE TEST <u>USE</u> HAVE	,SELTERR=NO <u>YES</u>
	,DYNTAPE	,MSGDGD= <u>DEFERRED</u> ACTIVE ROLLEDOFF INPUT
	,DYNTAPE2	,VRECAT

RESTORE STATEMENT The RESTORE TYPE=ABR statement requests a restore of individual data sets from Volume Backups. Only one RESTORE statement is allowed per execution of ABR, but any number of data sets can be restored in one ABR step.

SIMREST STATEMENT If SIMREST is coded, ABR will print the data sets which will be selected and the backup volumes necessary to do the restore. The data sets are not restored. It can be used to test RESTORE options or to pre-pull the tapes required.

FDRCLONE FDRCLONE is a separately-licensed enhancement to FDR/ABR. It is described in detail starting in Section 50.90.

FDRCLONE is a facility which "clones" disk volumes or data sets belonging to one MVS system to another MVS system, either an LPAR or a separate system. Its input is your normal FDRABR VOLUME backups, but data sets are actually restored only when they are needed. You can clone:

- All disk volumes in a data center
- selected disk volumes
- selected data sets

As cloned data sets are needed, they are dynamically restored from regular FDRABR volume backups (full-volume and incremental ABR backups) as described in this section. Only data sets which are actually needed by batch jobs or TSO users will be restored, so the total size of the restored data may be much less than the total in use at your home site.

50.08 CONTINUED . . .

**ABR DATA SET
RESTORE
PROCEDURE**

T his restore procedure restores data sets which were backed up by ABR Volume Backups. Usually these data sets will still exist on their original disk, so ABR will restore them in place, overlaying their existing contents. If they do not exist on disk or if a restore to a new name which does not exist is requested, ABR will allocate them.

[Section 50.02](#) details how ABR restore selects the appropriate ABR backup file on disk or tape for each data set to be restored. If multiple data sets are to be restored from a given backup, that file will be read only once. While reading a backup file, ABR can restore the selected data sets to one or more disk volumes concurrently. The target disk volume will be selected for each data set by the following rules:

- If the NVOL= operand was specified on the SELECT statement which selected this data set, that volume or volumes will be used. See the description of NVOL= in [Section 50.09](#) for details of target volume selection.
- If the output data set name is cataloged, then the volume to which it is cataloged will be chosen. The output data set name will be the original name or the new name if a NEWNAME=, NEWGROUP=, or NEWINDEX= operand was specified on the SELECT statement which selected the data set. If the data set is cataloged as being on multiple volume serials, then the volser will be selected from that list based on the volume sequence number in the F1 DSCB (field DS1VOLSQ) of the input data set.
- If none of the above apply, then the serial of the volume from which the data set was dumped, as recorded by ABR, will be used.
- If the data set was not preallocated on the selected volume, and the allocation fails on that volume for any reason, the ABR RESTORE ALLOCATION LIST, if enabled ([See Section 90](#)), will be checked to see if there is an ALLOCATE statement which applies to this data set. If so, the NVOL list from that statement will be used as described above for NVOL=. The RESTORE ALLOCATION LIST can be used to identify alternate volumes to which to restore data sets if their target volume is full or no longer exists.

If SMS (System Managed Storage) is active on this system, and the data set does not already exist on the volume selected by the rules above, SMS is invoked to decide if the data set should be SMS-managed. If so, SMS will select an output volume. SMS rules are detailed in [Sections 20.01](#) and [70](#).

**RESTORE ICF
VSAM FILES**

ABR will restore ICF VSAM files using the base cluster name. ABR will restore each individual component associated with this cluster name, and will allocate ICF VSAM files if they do not currently exist. ABR will update the appropriate fields within the VVR for each component. ICF VSAM files except the VVDS itself are movable; except for the VVDS and catalogs they can be restored to a new name or group. If NEWGROUP= or NEWINDEX= are specified the new group name will be applied to both the cluster name and all of its components. If NEWNAME= is specified for a cluster which is not allocated, ABR will let VSAM determine names for the components. Some information contained solely in the catalog, including protection (RACF or password) and expiration date will not be updated. However, path names for alternate indexes (AIXs) and aliases of user catalogs will be restored as long as both backup and restore are done with V5.3 level 30 or above. [See Section 80.13](#) "VSAM Special Considerations" for a more detailed explanation.

50.08 CONTINUED . . .

OPERANDS	TYPE=ABR	Required to restore data sets from ABR Volume Backups (regardless of whether the dump was done with TYPE=FDR, TYPE=ABR, TYPE=DSF, or TYPE=AUTO). ABR will attempt to restore all the data sets specified by the SELECT statements, DD=ALL option or remote queue data set. If the original data set does not exist on the volume and the data set is not recorded in the ABR Scratch Catalog, a manual restore (GEN= and CYCLE=, or TAPEDD=) must be specified. A TYPE=ABR restore using the TAPEDD= option can also be used to restore from any backup created by any part of the FDR system (FDR, DSF, or ABR).
	BLKF=	<p>Changes the blocksize (BLKSIZE) of certain fixed (RECFM=FB) and variable (RECFM=VB) data sets. BLKF= specifies the number of blocks per track (1 to 10) and is used to calculate a new blocksize based on the output track size (BLKF=1 results in BLKSIZE=32760). For RECFM=FB data sets the blocksize will be rounded down to a multiple of the logical record length (LRECL).</p> <p>DSORG=PS (sequential) data sets will be physically reblocked during the restore; the new blocksize must be larger than the original blocksize except when restoring from 3390 to 3380.</p> <p>DSORG=PO (partitioned) data sets will not be reblocked but the new blocksize will be used for new members; it must always be larger than the original blocksize.</p> <p>Other data set types, including UN (undefined), DA (direct access), PSE (enhanced format sequential), POE (PDSE), HFS and all types of VSAM, will not be reblocked.</p> <p>Default: data sets are not reblocked during restore; all original blocks will be restored without change, although they may be written to new locations. BLKF= is usually used when restoring to an unlike device type (e.g., 3380 to 3390) to improve utilization of the new track size but can also be used during like device restores.</p>
	BYPASSACS	<p>On a system with SMS (System Managed Storage) active, the SMS ACS (Automatic Class Selection) routines are not to be invoked for data sets which must be allocated. If a data set has a SMS storage class assigned (see STORCLAS= in Section 50.09) it will be SMS-managed, and SMS will be invoked to allocate the data set on an SMS-chosen volume, but SMS will not be allowed to override the storage class or management class assigned to the data set.</p> <p>Default: on an SMS system, the SMS ACS routines will be invoked for every data set which has to be allocated. The assigned storage and management classes will be passed to those routines, which can approve or override them. A data set will be passed to SMS for allocation if the storage class ACS routine assigns a storage class to the data set.</p>

50.08 CONTINUED . . .

BYPASSSMS On a system with SMS (System Managed Storage) active, SMS data sets will be directly allocated on SMS-managed volumes, bypassing normal SMS storage group and volume selection. The selected output volume must be a SMS-managed disk volume, and the data sets being restored must have a SMS storage class assigned (see BYPASSACS above and STORCLAS= in [Section 50.09](#)). The data sets will be allocated and cataloged according to SMS standards.

Normal SMS facilities do not allow allocation of data sets on specific volume serials, but BYPASSSMS will do so, allowing data sets to be located for performance or other reasons. Note that if BYPASSACS is also specified, the assigned SMS classes will not be validity- or authority-checked.

Default: on an SMS system, for data sets which are SMS-managed and must be allocated, the SMS storage group ACS routine will be invoked to select a storage group and SMS will select a SMS-managed volume and allocate and catalog the data sets.

BYPASSACS and BYPASSSMS are primarily for use by storage administration personnel, since they bypass normal SMS allocation controls and rules. In order to use BYPASSACS or BYPASSSMS, the user of ABR must be authorized to the RACF profile

STGADMIN.ADR.RESTORE.BYPASSACS

in class FACILITY, or the equivalent in other security systems.

CATIFALLOC Specifies that non-VSAM output data sets will be cataloged even if they were preallocated (not allocated by ABR); the output data set will be cataloged if it is not already cataloged on another volume (unless the RECAT operand was specified).

Default: output data sets are cataloged only when the restore allocates them.

COPY= Specifies the copy of the backup from which the restore is to be attempted; "n" can be any digit from 1 to 9. COPY=2 can be specified if a duplicate backup (TAPExx) created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utilities ([See Section 60](#)).

Default is COPY=1 unless overridden in the FDR Global Option Table ([See Section 90](#)).

DATA= **ALL** – all of the allocated tracks in each data set will be restored. ALL should not be specified unless the backup was taken using the DATA=ALL option. DATA=ALL should not be specified with the RLSE and %FREE operands.

USED – only the used portion of the PS (physical sequential) and PO (partitioned) data sets will be restored.

Default: USED.

DD=ALL Specifies that ABR is to scan all the DISKxxxx DD statements for the DSNNAME coded. ABR will attempt to restore these data sets, unless they are temporary data sets, processing them as though they had been requested on a SELECT statement that specified DSN= with no other operands. DSNENQ=NONE is the default for these data sets.

50.08 CONTINUED . . .

DSNENQ=

Specifies whether all of the data sets being restored will be ENQed. See “Data Set Enqueue Option” in [Section 50.03](#) for more details.

If you are restoring over an existing data set and the ENQ fails, the restore will be bypassed with an error message. If the restore must allocate the output data set and the ENQ fails, no error message is issued and the restore is still done. A successful ENQ will prevent any other task from using the data set until the restore from the current backup data set is complete. An ENQ failure is considered an error but it will not prevent other data sets from being restored. The options for DSNENQ= are:

USE – The data sets will be enqueued for the duration of the restore from the current backup data set. This is the most frequently used option.

TEST – The data sets will only be tested to see if they are enqueued to another task at the time the restore starts. The data set will not be enqueued and other tasks may enqueue it while the restore is preceding.

HAVE – The data sets will be enqueued for the duration of the restore. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not enqueue the data set)

RETRY (try the enqueue again; may result in the FDRW27 message again)

NONE – No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: USE. Note that NONE or TEST may allow other jobs to attempt to read the data set being restored before ABR has restored all of the data tracks.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the restore is complete. However, use DSNENQ=NONE when another data set by the same name on another volume may be in use (e.g., restoring data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

**DYNTAPE
DYNTAPE2**

Specifies that ABR is to dynamically allocate the backup data sets using a DDNAME of TAPE#. This option should be used if the backup is on disk, in a automated tape library (ATL or silo) or a mix of backup device types will be read, such as 3480 and 3490E. Do not use DYNTAPE and TAPEDD= in the same execution.

DYNTAPE2 will allocate 2 drives which will improve performance when restoring from multi-volume tape backups. DYNTAPE2 should not be used if backup files on disk might be involved.

Default: the first TAPEx DD statement found in the step JCL is used to mount all backups, unless TAPEDD= was specified.

50.08 CONTINUED . . .

ICFCAT=	<p>Applies to ICF VSAM files only. Specifies the source of the catalog name to be used if an output ICF VSAM cluster must be allocated.</p> <p>ORIGINAL – use the catalog in which the original dumped cluster was cataloged. When restoring a cluster to a new name, ICFCAT=ORIGINAL is treated like ICFCAT=ALIAS, described below. If you need to catalog the output cluster into the same catalog as the input cluster but that catalog is not the one aliased for the new name, you must specify ICFCAT=STEPCAT and supply a STEPCAT DD statement pointing to that catalog.</p> <p>STEPCAT – use the STEPCAT as the target catalog. If a STEPCAT DD statement is not supplied, it will use the master catalog or the catalog which is aliased for this data set in the master catalog.</p> <p>ALIAS – determine the catalog from the alias name in the master catalog. If no alias is found and the cluster is being restored to the same name, use the input cluster's original catalog. If no alias is found, and the cluster is being restored to a new name, it will use the STEPCAT (if present in the JCL) or the master catalog. Multi-level alias (MLA) is supported.</p> <p>Default: ORIGINAL, except that if the cluster is being restored to a newname (NEWGROUP or NEWINDEX specified) the default is ALIAS. If the output cluster is SMS-managed, ALIAS is forced.</p> <p>Note: if you receive message FDR157 with COMP=0004 and CODE=00120 for a cluster, this means that its original catalog name does not exist on this system. Use ICFCAT=ALIAS or STEPCAT to place the cluster into a catalog which exists.</p>
MAXCARDS=	<p>Accept additional SELECT and EXCLUDE statements (over 100).</p> <p>Default is 100 statements.</p>
NOCAT RECAT	<p>NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.</p> <p>RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.</p> <p>Default: catalog output non-VSAM data sets only if they are not currently cataloged.</p> <p>NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.</p> <p>NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified.</p>
OPERATOR	<p>Requests that an operator message (FDRW24) will be issued for each tape necessary to complete the restore. This option gives the operator the ability to pre-pull required tapes or bypass a tape which may not be available at this time.</p>

50.08 CONTINUED . . .

PRESTAGE	<p>Output data sets which already exist on the target output volume will not be restored. This may be used to avoid restoring data sets which have already been restored. If the output data sets do not exist on the target volume, they will be allocated and restored.</p> <p>Default: pre-allocated data sets will have their contents overlaid.</p>
RLSE %FREE=	<p>RLSE – all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn – a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement).</p>
SELTERR=	<p>Specifies what will happen at step termination if one or more of the SELECT or EXCLUDE statements was never referenced (no data set to be restored was selected by the statement):</p> <p>NO – a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.</p> <p>YES – a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.]</p> <p>Default: YES unless overridden in the FDR Global Option Table (See Section 90).</p>
SMSGDGD=	<p>Specifies the status of SMS-managed GDG (Generation Data Group) data sets, if allocated by the restore.</p> <p>DEFERRED, ACTIVE, or ROLLEDOFF will set the GDG to that status.</p> <p>INPUT will set the GDG to the original SMS status of the GDG generation, as recorded on the backup tape. If the original GDG was non-SMS, it will be set ACTIVE if that generation is currently cataloged, otherwise DEFERRED.</p> <p>If a GDG is restored as DEFERRED, you may need to execute an IDCAMS "ALTER ROLLIN" to make the generation active.</p> <p>Default: DEFERRED.</p>

50.08 CONTINUED . . .

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume).

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for some reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

50.09 SELECT STATEMENT FOR DATA SET RESTORE

SELECT	DSN=filter	,NOCAT
S	CATDSN=filter	,RECAT
	DD=ddname	
EXCLUDE	ALLDSN	,NOTIFY=userid
X		
	,BLKF=nn	,NVOL=(vvvvvv,vvvvvv,...)
	,CATALOG=catname	,OLDBACKUP=nn
	,MCATALOG=catname	
	,CATLIMITGDG=n	,PRESTAGE
	,COPY=n	,PRTALIAS
	,CYCLE=nn	,RLSE
		,%FREE=nn
	,DATA=ALL NONE	,STORCLAS=storageclass
	,DATACLAS=dataclass	,NULLSTORCLAS
	,NULLDATACLAS	,TAPEDD=x
	,DSNENQ=NONE	,TRK=nnnnn
		,CYL=nnnnn
	,MGMTCLAS=managementclass	
	,NULLMGMTCLAS	,VOL=vvvvvv
	,GEN=nnnn	,VRECAT
	,NEWNAME=newdsname	
	,NEWGROUP=newgroup	
	,NEWINDEX=new index	
	,NEWDD=ddname	

**SELECT
STATEMENT**

This statement is used with RESTORE TYPE=ABR and selects the data sets to be restored from full-volume or incremental backups. The **SELECT** statement identifies the individual data set name or group of data sets to be processed. The **EXCLUDE** statement identifies data sets from within those selected by SELECT statements which are not to be processed. As described in [Section 50.02](#), ABR will select the backups of the selected data sets using information in the VTOC of the data set's volume or the ABR scratch catalog, locate the indicated backup on tape or disk, and restore it. EXCLUDE statements should only contain the operands DSN=, DD=, ALLDSN, VOL=, or TAPEDD=. A maximum of 100 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since ABR will only restore data sets which are selected, EXCLUDE statements are required only to exclude certain data sets from within a larger group on a SELECT statement.

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A . B . **
SELECT  DSN=A . **
```

Example 2: Select all data sets on the backup of volume TSO001 (generation 27, cycle 00) except those beginning with "ABC":

```
EXCLUDE DSN=ABC**
SELECT  ALLDSN, VOL=TSO001, GEN=27, CYCLE=00
```

50.09 CONTINUED . . .

**NEWNAME/
NEWGROUP/
NEWINDEX for
ICF VSAM**

If you are restoring to a pre-allocated ICF VSAM cluster, where the cluster name is the same as the original, but the components may be named, you must specify the cluster name as both DSN= and NEWNAME=; this causes ABR to LOCATE the new component names. NEWNAME= cannot be specified for clusters with more than one alternate index.

If restoring an ICF VSAM cluster to a new name, if the new cluster must be allocated, you should specify NEWGROUP= or NEWINDEX=. ABR will modify both the cluster and component names.

VSAM catalogs and VVDSs cannot be restored to a NEWNAME.

See [Section 80.13](#) for "VSAM Special Considerations".

OPERANDS DSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the names of data sets on the backup tapes to be restored. For ICF VSAM clusters, only the cluster name is compared; you cannot select by component name.

```
EXAMPLES : DSN=USER1.JCL.CNTL
           DSN=**LIST
           DSN=PROD++.**.LIB*
```

You cannot use DSN= to select generation data group (GDG) generations by relative generation number. Use CATDSN= or DD= if relative GDG numbers are required.

If DSN= specifies a fully-qualified name, ABR will attempt to locate the requested backup of that data set as described in [Section 50.02](#), and restore it.

If DSN= specifies a filter (any wildcard characters), one specific backup file must be identified, either by the VOL=, GEN= and CYCLE= operands, or the TAPEDD= operand; ABR will read only that specific backup and will not search for the other backups.

50.09 CONTINUED . . .

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= parameter, causing ABR to look on those volumes for the DSCB containing backup information). Specification of a relative generation number for GDG data sets is supported, e.g., CATDSN=A.B(-1)

If a filter is specified, then catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT DSN=dsname, VOL=volser was present for each of them. It may be necessary to specify MAXCARDS=nnnnn if a large number of data sets are selected by the filter.

Additional considerations for CATDSN=filter are explained in [Section 80.14](#).

CATDSN= is supported only on SELECT statements. However, a preceding EXCLUDE statement with DSN= and/or VOL= can exclude data sets from selection by CATDSN=.

If the VOL= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

```
EXAMPLES: CATDSN=USER1.JCL.CNTL
          CATDSN=**MASTER(0)
          CATDSN=PROD++.**.LIB*
```

Normally CATDSN= will not display the data sets it selects from the catalogs, you will see the names only when FDRCOPY actually finds and selects the data sets in the VTOCs of the volumes they are cataloged to. To display all of the data sets selected specify PCATDSN=filter.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

DD=

Specifies that a data set name is to be taken from a DD statement. This operand must point to the DDNAME of a JCL statement. Using this option enables you to specify a non-standard data set name or a generation data set (GDG) relative generation.

```
EXAMPLE:  SELECT      DD=DD1
           //DD1      DD   DSN=A.B.C(0),DISP=SHR
```

ALLDSN

All data sets found on a specified backup file are to be restored. To use ALLDSN one specific backup file must be identified, either by the VOL=, GEN= and CYCLE= operands, or the TAPEDD= operand; ABR will read only that specific backup and will not search for the other backups.

NOTE: DSN=, CATDSN=, DD= and ALLDSN are mutually exclusive. One and only one must be specified.

50.09 CONTINUED . . .

BLKF=	<p>Selected PS and PO data sets are to be reblocked during the restore; see BLKF= in Section 50.08 for details.</p> <p>Default: data sets are not reblocked unless BLKF= was specified on the RESTORE statement. The restore will fail if the input data set has blocks larger than the track size of the output disk.</p>
CATALOG= MCATALOG=	<p>Specifies the name of a user catalog (CATALOG=) or alternate master catalog (MCATALOG=) to search when CATDSN= is specified. See Section 80.14 for details.</p> <p>Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter.</p>
CATLIMITGDG=	<p>May be used with CATDSN=filter to limit the selection of GDGs from the catalogs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:</p> <p>n will cause only the most recently created "n" generations to be selected. -n will cause only generation (-n) to be selected.</p> <p>Default is that all the generations of selected GDGs will be selected (unless a relative generation number is specified at the end of the filter, e.g., CATDSN=filter(-2)).</p>
COPY=	<p>Specifies the copy of the backup from which the restore is to be attempted; "n" can be any digit from 1 to 9. COPY=2 can be specified if a duplicate backup (TAPExx) created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utilities (See Section 60).</p> <p>Default is the value from the RESTORE statement or from the FDR Global Option Table (See Section 90).</p>
CYCLE=	<p>Specifies the cycle number of the backup ABR is to use to restore this data set. This cycle number is the last two digits of the backup data set name and is displayed in various ABR reports. Use this option if a data set was scratched and was not recorded in the ABR SCRATCH Catalog, or if you wish to restore from other than the most current backup. It's value can range from zero to 63. CYCLE= must be specified with GEN= and VOL=.</p> <p>If OLDBACKUP was enabled on the data set's disk volume when dumped, the OLDBACKUP= operand can also be used to request an older backup.</p>

50.09 CONTINUED . . .

DATA=	<p>ALL – restore all allocated tracks in each data set selected by this SELECT statement.</p> <p>NONE –no data tracks will be restored for any data set type. However, if output data sets need to be allocated, the restore will use the control information on the backup data set to allocate the data set; also, the characteristics of the data set will be updated from the backup even if the data set was preallocated. This allows you to preallocate output data sets with a restore without restoring data and recover their contents from another source (such as a data base backup). You must use DATA=NONE if the backup was done with DATA=NONE (see Section 50.06), otherwise you will get error messages indicating data tracks were missing.</p> <p>Default: restore only the used tracks of PS and PO data sets, unless DATA= was specified on the RESTORE statement.</p>
DATACLAS= NULLDATACLAS	<p>On a system with SMS active, specifies the SMS data class to be associated with the data set being restored, overriding the original data class of the data set (if any). The Data Class ACS routine will not be invoked.</p> <p>NULLDATACLAS changes the data class to null (not specified).</p> <p>Default: the original data class of the input data set (if any) will be associated with the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is set.</p>
DSNENQ=	<p>NONE – bypass the data set enqueue for data sets selected by this SELECT statement.</p> <p>Default: the enqueue option is determined by the DSNENQ= operand specified on the RESTORE statement.</p>
GEN=	<p>Specifies the generation of the backup ABR is to use. If GEN= is used, CYCLE= must also be specified. This value is the 'gggg' portion of the tape data set name, which has the Format FDRABR.Vvvvvvv.Cnggggcc. This option is used to restore from a previous generation of backups or if the information ABR recorded is incorrect. Specifies a value from 1 to 9999.</p> <p>Default: ABR will locate the most recent backup of the data set and restore it.</p>
MGMTCLAS= NULLMGMTCLAS	<p>On a system with SMS active, specifies the SMS management class to be presented to the SMS Management Class ACS routine for the data set being restored, overriding the original management class of the data set (if any). The ACS routine may accept or override this class.</p> <p>NULLMGMTCLAS changes the management class to null (not specified).</p> <p>Default: the original management class of the input data set (if any) will be passed to the ACS routine for the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is passed.</p>

50.09 CONTINUED . . .

NEWNAME=
NEWN= Restore the selected data set with a new name. NEWNAME should only be used with DSN= or DD=, and should not be used for ICF VSAM clusters unless they are preallocated. If the newname ends in a GDG relative generation number, e.g., NEWNAME=gdgname(-1), a LOCATE will be done to get the proper absolute generation number.

NEWGROUP=
NEWG= Restore the selected data sets using a new group name. The characters specified will replace the beginning of the input data set name. Care should be taken when periods are used that index levels are not incorrectly changed. ABR will check the new names for valid IBM standards. For ICF VSAM, the new group is applied to the cluster name and all component names.

EXAMPLE: SELECT DSN=ABC**, NEWG=XYZ

Any data sets restored will be renamed to start with characters XYZ.

50.09 CONTINUED . . .

**NEWINDEX=
NEWI=**

Restore the selected data sets using a new name formed by adding or replacing one or more index levels in the original name; replacement index levels do not have to be the same length as the original indexes they replace. In the simplest case, each index level specified in NEWI is used in place of the corresponding index in the original name. Any remaining index levels at the end of the name are copied unchanged. This can easily be used to change the first indexes of the name.

For example, if the input data set is A.B.C.D,

NEWI=D	results in	D.B.C.D	(first index replaced)
NEWI=DD.E	results in	DD.E.C.D	(first 2 indexes replaced)

If a period is specified without any preceding characters, one original index level is copied from the input data set name to the output. This allows you to easily modify indexes in the middle of the name.

For example, if the input data set is A.B.C.D,

NEWI=..E	results in	A.B.E.D	(third index replaced)
NEWI=FF...G	results in	FF.B.C.G	(first and fourth indexes replaced)

If + is specified before a new index level, that new index is inserted into the output data set name at that point. If ++ precedes the new index, it will be added to the end of the name. If - is specified, the next input index level will be dropped from (not copied to) the output name.

For example, if the input data set is A.B.C.D,

NEWI=+F	results in	F.A.B.C.D	(new first index added)
NEWI=..+F	results in	A.B.F.C.D	(new third index added)
NEWI=++F	results in	A.B.C.D.F	(new last index added)
NEWI=.-	results in	A.B.D	(third index dropped)
NEWI=Q.-.+E	results in	Q.C.E.D	(combination)

Note that, except for the ++ option, every period in the NEWI= mask corresponds to a period (one index level) in the original (input) data set name. The resulting new name will be checked to insure it meets IBM standards.

If the NEWI= value ends in a GDG relative generation number, e.g., NEWI=..NEWMAST(-2), that relative number will be added to the end of the newname, and a LOCATE done to get the proper absolute generation number.

NEWI= is a convenient way to rename every input data set, while using some index levels from the original name and replacing other indexes or adding new indexes. . For ICF VSAM, the NEWINDEX is applied to the cluster name and all component names.

NEWDD=

Specifies the name of a DD statement from which the new name of the output data set is obtained.

NOTE: NEWN=, NEWG=, NEWI=, and NEWDD= are mutually exclusive. If none of them are specified, the data set is restored under its original name. NEWN= and NEWDD= should not be used on SELECT statements which select more than one data set.

50.09 CONTINUED . . .

**NOCAT
RECAT**

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged, unless overridden by NOCAT/RECAT on the RESTORE statement.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.

NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified on the RESTORE statement.

NOTIFY=

Specifies that ABR is to notify the TSO Userid specified at the completion of the restore. The userid is from 1 to 7 characters.

50.09 CONTINUED . . .

NVOL=

Specifies the volume serial(s) of output disk volumes to which data sets selected by this statement are to be restored. You may specify:

- 1) A single specific volume serial, e.g.,
NVOL=ABC123
- 2) A list of specific volume serials, enclosed in parentheses, e.g.,
NVOL=(TSO001,TSO002,TSO003)
- 3) A group of volumes by placing an asterisk at the end of the volsr prefix, e.g.,
NVOL=TSO*
- 4) A combination of specific and group, e.g.,
NVOL=(TSO*,PROD*,ABC001)
- 5) All online disk volumes may be selected by
NVOL=*

A list of online target volumes matching your specification is generated by scanning all disk UCBs in the system UCB chains; there is no guarantee of the order in which UCBs are found, so you cannot predict the order of the volume serials in the list. If you specify volume serials or groups which are not online, they are ignored and no error message will result.

However, if the first or only specification is a specific volume serial, it will be chosen as the first target volume, with other volumes placed after it in UCB chain order.

Also, if you are restoring a multi-volume data set (non-VSAM or SMS-managed VSAM), the volume sequence number of the piece of the data set being allocated will be used to select a specification from your list. For example, if NVOL=(A,B,C), the second piece of the data set will go to volume B. If that specification is a group, the first volume in the UCB chain matching that group will be tried. If the allocation is unsuccessful (such as insufficient free space), then other volumes in the NVOL list will be tried as described above.

The first target volume is checked to see if an output data set already exists there. If so, it restores over the existing allocation (unless PRESTAGE was specified). If not, it attempts to allocate the output data set on that volume. If the allocation fails, it will be retried on successive volumes in the list until it succeeds or until it fails on 64 volumes. If the list contains several disk device types, "like" volumes (same type as the data set being restored) will be tried first, then unlike devices.

For multi-volume data sets, a target volume is bypassed if a piece of the data set already exists there but is not the right piece, so that it will not attempt to restore the third volume of a data set on top of the first volume. When it finds a target volume in the list that does not contain a piece of the data set, it will be allocated.

Specifying multiple volsers or a volume group allows you to restore data sets in one pass even when no one volume has available space to contain them all; they will be spread across many of the target volumes.

Default: the output volume will be selected by rules defined in [Section 50.08](#). Note that when NVOL= is specified, and data sets are selected which are currently allocated and cataloged, ABR will restore them to the new volumes, and not to the volume on which they are cataloged.

On a system with SMS active, NVOL= may be ignored if the data set does not exist on the volume specified and the data set is SMS-managed (see STORCLAS= below).

Note: if an allocation is attempted on several volumes from your NVOL list but it fails on all of them, the message printed will usually show the allocation failure codes from the **first** volume only; failure codes from other volumes are not displayed and may be different.

50.09 CONTINUED . . .

OLDBACKUP=	Specifies that ABR is to restore from a backup of this data set taken nn backups ago. OLDBACKUP must be enabled in the ABR Model DSCB on the volume involved. nn is from 0 to the maximum number of old backups recorded for the data set. "1" specifies the backup taken prior to the current backup for this data set, "2" the next oldest, etc. If not specified (or 0) the most current backup will be restored unless GEN= and CYCLE= or TAPEDD= are specified. See Section 50.02 for further explanation.
PRESTAGE	Specifies that selected data sets will not be restored if the output data set already exists on the first target output volume. This may be used to avoid restoring data sets which have already been restored. Default: restore pre-allocated data sets, overlaying the existing contents of those data sets, unless PRESTAGE was specified on the RESTORE statement.
PRTALIAS	When used on a SELECT statement with CATDSN= it will display all of the alias names and user catalogs which were searched.
RLSE %FREE=	RLSE – specifies that all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released. %FREE=nn – specifies a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space. Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation. Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement or by RLSE/%FREE= on the RESTORE statement).
STORCLAS= NULLSTORCLAS	On a system with SMS active, specifies the SMS storage class to be presented to the SMS Storage Class ACS routine for the data set being restored, overriding the original storage class of the data set (if any). The ACS routine may accept or override this class. NULLSTORCLAS changes the storage class to null (not specified). Default: the original storage class of the input data set (if any) will be passed to the ACS routine for the output data set. For a non-SMS input data set, a null class is passed. If the Storage Class ACS routine assigns a storage class to this data set or accepts the class passed, the data set will be allocated as SMS-managed, and the SMS Storage Group ACS routine may be invoked to determine the actual target volume. If the Storage Class ACS routine returns a null (blank) storage class name, the data set will be allocated as non-SMS and the ABR rules listed in Section 50.08 for volume selection must select a non-SMS target volume.

50.09 CONTINUED . . .

TAPEDD=

Specifies the same character as specified in a TAPEX DD statement. This option is used to override ABR's selection of the backup tape from which to restore this data set. A TAPEX DD statement must be provided giving all the values necessary to point to the backup file to be restored from.

```
EXAMPLE :
//TAPEA DD DSN=FDRABR.VSCRO10.C1000101,DISP=OLD,LABEL=(2,SL),
//          UNIT=(TAPE,,DEFER),VOL=SER=000150
          RESTORE TYPE=ABR
          SELECT   DSN=ABC,TAPEDD=A
```

NOTE: If you provide more than one TAPEX DD in a step, UNIT=AFF=TAPEX should be used to reduce the number of tape drives required. Also, if normal restore is requested in the same execution, the first tape DD statement is used by ABR for these restores. TAPEDD must not specify this tape.

Do not use DYN TAPE and TAPEDD in the same execution.

**TRK=
CYL=**

If the data set selected by this SELECT statement must be allocated, CYL= or TRK= specifies the number of cylinders or tracks to be allocated to the data set. On PS or PO files when DATA=ALL was not specified, this value should be at least equal to the used portion of the data set. On all other types of files and when DATA=ALL is specified, this value should be equal to or greater than the original size of the file. For ICF VSAM clusters, modifies the size of the base data component only. If the space is too small for the data being restored, the restore will automatically extend the file for non-VSAM.

Default: use the original size of the data set.

VOL=

Specifies the disk volume serial number from which the data set name was dumped. VOL= must be specified if the original data set is not cataloged or recorded in the Scratch Catalog.

(See [Section 50.02](#) for a more detailed explanation).

50.09 CONTINUED . . .

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume) unless VRECAT was specified on the RESTORE statement.

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for any reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

50.10 FULL VOLUME RESTORE STATEMENT

```

RESTORE TYPE=FDR ,EXPDT=yyddd|yyyyddd
R      ,CONFMESS=YES|NO ,MAXERR=nn
      ,COPY=n ,ONLINE
      ,CPYVOLID=NO|YES ,OPERATOR
      ,DYNTAPE ,SMSPROT=ALL|NONE
      ,DYNTAPE2 ,VOLRESET=YES|NO

```

RESTORE STATEMENT The RESTORE TYPE=FDR statement performs a full-volume recovery of one or more DASD volumes from ABR Volume Backups. Only one RESTORE statement is allowed per execution of ABR.

ABR FULL VOLUME RECOVERY PROCEDURE ABR full-volume recovery reads a generation of Volume Backups, starting with an incremental backup (from which the VTOC, VTOCIX, VVDS and label track are restored, plus any data sets backed up to that incremental). It then reads successive incrementals in descending cycle order, restoring the most recent backup of each data track. Finally it reads the full-volume backup which began the generation, restoring all data tracks which were not found on any of the incrementals. The result is that the restored volume will look **exactly** as the volume did at the point of the first incremental backup read.

By default, ABR will start with the most recent incremental backup in the current generation, but if you want to restore to a different point-in-time, you can specify the cycle (incremental) to start with. You can even restore from an earlier generation. If the most recent incremental is a TYPE=DSF manual backup, it will be automatically bypassed since it does not contain all the information required, such as VTOC, VTOCIX, VVDS and label track.

If you are performing this restore at a disaster/recovery site, you must do a separate restore of the ABR catalog, since volume recovery depends on the current ABR catalog to get backup tape information. Although the ABR catalog will be backed up as part of daily incremental backups, you cannot depend on this backup since it will be done in the middle of ABR backups and will not contain the latest information for all volumes. [See Section 80.20 "Disaster Recovery"](#) for suggestions on the procedures to follow.

CAUTION: If data sets have been excluded from incremental backups, they will not be restored to the most current level, unless they have been unchanged since the full-volume backup at the beginning of the generation. This is true if the data set was excluded using the protect list or EXCLUDE statement

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data..

SIMULATED FULL VOLUME RECOVERY SIMREST is not supported with TYPE=FDR. If you wish to see which backup volumes FDR will need to restore a given disk volume, use program FDRABRP with the statement PRINT CATLG,VOL=vvvvv ([See Section 53](#)).

50.10 CONTINUED . . .

FDRCLONE AND FDRDRP ABR full-volume restores can be enhanced if you are licensed for FDRCLONE, documented starting in [Section 50.70](#). FDRCLONE is particularly important if you are backing up to high-capacity tape cartridges, such as the IBM Magstar or StorageTek Redwood or 9840.

FDRCLONE allows you to restore only the data sets that are required at a disaster/recovery site or a test system, instead of restoring all data sets including those which will never be used on that system. FDRCLONE restores data sets from ABR backups on demand, as they are needed.

FDRDRP is part of FDRCLONE. It is a full-volume restore program, with the same result as an ABR full-volume restore, but it manages the backup tapes to minimize tape mounts and reduce the elapsed time of the restores. Normal ABR full-volume restore does one disk volume at a time, and may mount an input tape multiple times. FDRDRP does many disks in parallel and mounts input tapes a minimum number of times (often only once).

OPERANDS	TYPE=FDR	Required for a full volume restore. It is followed by one or more SELECT statements, each specifying a volume to be recovered.
	CONFMESS=	<p>YES – before beginning the restore, FDR will request confirmation via a WTOR (FDRW01) message to which the MVS operator must reply.</p> <p>NO – suppresses the WTOR and begins the restore immediately.</p> <p>Default: YES.</p> <p>NOTE: CONFMESS=NO can be very useful at a disaster recovery site to avoid delaying full volume restores while waiting for an operator response.</p>
	COPY=	<p>Specifies the copy of the backup from which the restore is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utility (See Section 60).</p> <p>If COPY=1 or 2 and ABR finds that one of the backup tapes is not cataloged under the copy specified, ABR will check to see if the other copy was created. If cataloged, ABR will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) ABR will automatically use the other copy (1 or 2) if it still exists.</p> <p>Default is COPY=1 unless overridden in the FDR Global Option table (See Section 90).</p>

50.10 CONTINUED . . .

CPYVOLID=

Specifies whether the volume serial number of the disk that was backed up will be restored, if the existing volser of the output disk is different (if the serials are the same, CPYVOLID is ignored).

YES – volume serial number of the output volume will be replaced with the original volume serial number of the disk which was dumped. If another online volume has the same serial, the restored volume will be placed offline at the end of the restore.

NO – the volume serial number of the output volume will be retained. See the VOLRESET= operand below.

Default: NO – unless the volume being restored was SMS-managed, when YES is forced.

Note: although full-volume ABR restore does not catalog data sets, any data sets which were cataloged to the original volume are automatically cataloged to the new volume when restoring with CPYVOLID=YES, assuming that you restore the system catalogs to the same point-in-time. If you use CPYVOLID=NO and do not later relabel the volume, data sets may need to be manually recataloged.

VSAM/SMS WARNING: See the notes under VOLRESET=NO below. CPYVOLID=YES is recommended for any volume containing a VVDS.

**DYNTAPE
DYNTAPE2**

Specifies that ABR is to dynamically allocate the backup data sets using a DDNAME of TAPE#. This option should be used if the backup is on disk, in a automated tape library (ATL or silo) or a mix of backup device types will be read, such as 3480 and 3490E. Do not use DYNTAPE and TAPEDD= in the same execution.

DYNTAPE2 will allocate 2 drives which will improve performance when restoring from multi-volume tape backups, this is especially important for 3480 tape cartridges. DYNTAPE2 should not be used if backup files on disk might be involved.

DYNTAPE and DYNTAPE2 will deallocate the tape drive if the restore of the next disk in this ABR step does not need the same tape volume. If the same tape volume is required for the restore of the next disk volume, it will remain mounted.

Default: the first TAPE# DD statement found in the step JCL is used to mount all backups, unless TAPEDD= was specified.

50.10 CONTINUED . . .

EXPDT=

Specifies an expiration date which will be passed to OPEN when each backup file is opened. Since these are input tapes, the expiration date will probably be ignored by your tape management system except for certain special dates. The most common use will be EXPDT=98000, which is accepted by most tape management systems and means "this tape is not in the tape management data base". You might need to use EXPDT=98000 (or whatever your TMS supports) when the tape management data base has not been restored to a point after the ABR backups were taken, so that it doesn't reflect the backup tapes that FDRDRP needs to read. In this case, it might be easier to disable tape management until the data base can be made current. If you use a TAPEX DD statement (instead of DYN TAPE), you can also specify the EXPDT= operand on the DD statement; the value on the DD will override this operand if both are specified.

The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

MAXERR=

Specifies the number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP. If the backup is written to the ExHPDM subsystem, MAXERR=1 is recommended.

Default: 20 errors.

ONLINE

Allows ABR full-volume restore to restore to online disk volumes without requiring DISKxxxx DD statements for them. The required output volumes (identified by VOL= or NVOL= operands) will be dynamically allocated when needed.

OPERATOR

Specifies that, before the RESTORE operation began, an operator message will be issued for each tape necessary to complete the restore. This option gives the operator the ability to pre-pull required tapes or bypass individual cycles for which the tapes may not be available at this time. However, if some cycles are bypassed, the restore may correctly restore the latest version of some data sets.

SMSPROT=

ALL – enforces several rules when SMS-managed volumes are involved: Backups of SMS-managed volumes can only be restored to SMS-managed volumes, and non-SMS volumes only to non-SMS volumes. CPYVOLID=YES is forced when an SMS-managed volume is restored.

NONE – allows the restore of SMS-managed volumes to non-SMS volumes, and vice versa. Also allows the restore of SMS volumes to new volsers if CPYVOLID=NO is specified.

WARNING: SMSPROT=NONE should be used with caution. It will usually be used at a disaster recovery site where a relPL of MVS will be done after the restores, to place all volumes in the proper SMS status. See Section 70 for more details on restoring and moving SMS-managed volumes.

Default: ALL.

50.10 CONTINUED . . .

VOLRESET=

When CPYVOLID=NO is specified or defaulted, and the volume serial of the output disk is different from that of the original disk on the backup data set, the backup is restored but the volume serial of the output volume is retained. VOLRESET= controls additional processing relating to this change in volume serials. VOLRESET is ignored if the volume serial of the output disk is not being changed.

YES – the volume serials that are part of the data set names of the VTOC Index ("SYS1.VTOCIX.volser") and the ABR Model DSCB (usually "FDRABR.Vvolser") are checked to see if they match the input volume serial (the volume backed up). If so, they are renamed to match the volser of the output disk. Also, the DSCB field DS1DSSN (data set serial number, usually the volser of the first or only volume of the data set) for every data set on the volume will be changed to the new volume serial if the existing value matched the original volume serial. In the ABR Model DSCB, the current cycle number is set to the cycle number of the first incremental read during the restore, so that the next incremental will produce the next cycle in this generation (a restore using TAPEDD= will set the cycle number to 63 to force the next backup to be a full-volume backup). VOLRESET=YES should be used with CPYVOLID=NO when you intend to permanently retain the volume serial of the output disk. Note that data sets on the volume will not be recataloged to the new volume serial.

NO – do not rename the VTOCIX and ABR Model DSCB, do not change DS1DSSN fields, and do not reset the cycle number in the ABR Model DSCB. VOLRESET=NO should be used with CPYVOLID=NO if you plan to eventually relabel the disk back to the original volume serial.

VSAM/SMS WARNING: ABR will not rename the VVDS, since the VVDS and the catalogs contain self-defining records that would also need resetting. If a volume containing VSAM clusters or SMS-managed data sets is restored with CPYVOLID=NO, the data sets will be inaccessible unless the volume is relabeled to the original serial number. CPYVOLID=YES is recommended for any volume containing a VVDS, but if it is necessary to restore the volume under a temporary volser, use CPYVOLID=NO and VOLRESET=NO and relabel the volume with ICKDSF later.

Default: YES but it is ignored if CPYVOLID=YES is also specified.

50.11 SELECT STATEMENT FOR FULL VOLUME RESTORE

```

SELECT      VOL=vvvvvv          ,GEN=nnnnn|CURRENT
S
              ,COPY=n            ,NVOL=vvvvvv
              ,CYCLE=nn

```

SELECT VOLUME STATEMENT One or more SELECT Volume statement must follow a RESTORE TYPE=FDR statement to specify the volumes to be recovered, one per statement. Volumes will be recovered one at a time, in the order specified.

OPERANDS VOL= Specifies the disk volume serial number of the original disk to be restored. It must be present. If NVOL= is not specified, the backup will be restored to a disk volume which currently has the original serial (either the original disk or a new volume which has been initialized with that serial). A DISKxxxx DD statement must be provided pointing to the restore target volume unless the ONLINE operand was specified on the RESTORE statement.

COPY= Specifies the copy of the backup from which the restore is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRtsel utility ([See Section 60](#)).

If COPY=1 or 2 and ABR finds that one of the backup tapes is not cataloged under the copy specified, ABR will check to see if the other copy was created. If cataloged, ABR will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) ABR will automatically use the other copy (1 or 2) if it still exists.

Default is the value from the RESTORE statement or from the FDR Global Option table ([See Section 90](#)).

CYCLE= Specifies the cycle (incremental backup) number which ABR is to read first while restoring this volume. This option is used if you do not wish to start the restore with the most recent cycle created in the current generation or if you specified GEN=gggg to restore from an older generation (CYCLE= is required if GEN= is given). ABR will read this backup file first and continue backwards through the incrementals until the full volume (TYPE=FDR) backup is encountered. CYCLE=0 will only restore the full volume backup.

Default: see GEN= below.

50.11 CONTINUED . . .

GEN= Specifies the generation number which ABR is to read while restoring this volume. May be used if you wish to restore from a specific generation other than the most current. If GEN= is specified, CYCLE= must also be specified.

If GEN=CURRENT is specified, ABR will search for the most recently created generation and cycle cataloged for this volume in the ABR catalog; the ABR Model DSCB is not used and original volume need not be online. If

GEN=CURRENT,CYCLE=nn is specified, then the specified cycle in the generation from the catalog will be used (CYCLE=00 can be used to restore only the full-volume backup from the current generation).

By default, ABR will read the ABR Model DSCB from the volume specified by VOL= to obtain the most recently created generation and cycle. If the volume is offline, if the model does not exist or has been destroyed, or if the model indicates that no backups exist, then it will default to GEN=CURRENT and get the GEN/CYCLE from the ABR catalog. So, if you want to restore from the latest incremental backup in the current generation, GEN= should be omitted.

NVOL= Specifies the current volume serial of the disk volume to which this backup will be restored. A DISKxxxx DD statement must be provided pointing to the restore target volume unless the ONLINE operand was specified on the RESTORE statement. This might be used at a disaster/recovery site where the site personnel pre-initialize all of the disks to known volume serials; you can do the ABR volume recoveries without having to relabel all of their volumes to your volsers. If CPYVOLID=YES was specified or defaulted on the RESTORE statement, the output volume will be relabeled to the original volume serial at the end of the restore.

Default: the backup will be restored to a disk labeled with the original volume serial

50.20 VOLUME BACKUP EXAMPLES

**FULL VOLUME
DUMP ALL
VOLUMES**

Take full-volume compressed backups of all disk volumes at the installation, except WORK volumes. Three TAPEX DD statements are provided, so three disk volumes will be concurrently dumped. VOL=(,,255) is specified in case any individual disk volume requires over 5 tapes. DSNENQ=USE causes ABR to print a FDR158 message for any data set that was in use at the time of the backup, but will back up the data set anyway (no abend will occur if data sets are in use). As each tape drive becomes available at the end of backing up a volume, ABR will piggyback the next disk volume as the next file on that tape.

NOTE: omit COMPRESS= if backing up to tape drives attached by ESCON or FICON channels.

```
//BACKUP      EXEC  PGM=FDRABR,REGION=6144K
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//SYSPRIN3    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//TAPE1       DD    UNIT=3480,DSN=FDR1,DISP=(,KEEP),VOL=(,,255)
//TAPE2       DD    UNIT=3480,DSN=FDR2,DISP=(,KEEP),VOL=(,,255)
//TAPE3       DD    UNIT=3480,DSN=FDR3,DISP=(,KEEP),VOL=(,,255)
//SYSIN       DD    *
DUMP          TYPE=FDR,ONLINE,DSNENQ=USE,ENQERR=NO,COMPRESS=ALL
EXCLUDE       ALLDSN,VOLG=WORK
```

NOTE: The dsnames on the TAPEX DD statements are dummy names supplied to satisfy the rules of JCL. ABR will generate the data set names for the backup files.

**FULL VOLUME
DUMP
CERTAIN
VOLUMES**

Take full-volume backups of all SYS and PROD volumes, dumping two volumes concurrently, and making two copies of each backup. The TAPE11 and TAPE22 DD statements will contain COPY 2 backups identical to the COPY 1 backups on TAPE1 and TAPE2. Two tape drives are assigned to each TAPE DD statement to minimize tape rewind wait time (which is an issue mainly on 3480 cartridge drives); VOL=(,,255) is required to cause MVS to pre-mount output tapes on the idle drives in each pair. The VTOCs of each disk volume will be ENQed and reserved during the backups. ABR will only dump the used portion of PS or PO data sets.

```
//BACKUP      EXEC  PGM=FDRABR,REGION=2048K
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//TAPE1       DD    UNIT=(TAPE,2),DISP=(,KEEP),VOL=(,,255),DSN=FDR1
//TAPE11      DD    UNIT=(TAPE,2),DISP=(,KEEP),VOL=(,,255),DSN=FDR11
//TAPE2       DD    UNIT=(TAPE,2),DISP=(,KEEP),VOL=(,,255),DSN=FDR2
//TAPE22      DD    UNIT=(TAPE,2),DISP=(,KEEP),VOL=(,,255),DSN=FDR22
//SYSIN       DD    *
DUMP          TYPE=FDR,ENQ=RESERVE,DATA=USED
MOUNT         VOLG=SYS
MOUNT         VOLG=PROD
```

50.20 CONTINUED . . .

INCREMENTAL BACKUP ALL ONLINE VOLUMES An incremental backup of all online disk volumes is done to two tape drives. Two disk volumes will be processed concurrently, creating multiple backup data sets on the tapes (one per disk volume). All data sets on the disks which have been updated since the last ABR full-volume or incremental backup of that volume will be selected. However, all data sets with the text "PAYROLL" anywhere in their name will be backed up regardless of their update status. All data sets selected will be printed in ABR VTOC format. The ABRBKDQ points to the ABR backup dump remote queue data set; if present, any queued requests for backup will be processed.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     UNIT=3590-1,DSN=FDR1,DISP=(,KEEP),VOL=(,,,255)
//TAPE2       DD     UNIT=3590-1,DSN=FDR2,DISP=(,KEEP),VOL=(,,,255)
//SYSIN       DD     *
              DUMP    TYPE=ABR,ONLINE,PRINT=ABR
              SELECT  DSN=**PAYROLL**
//ABRBKDQ     DD     DSN=FDRABR.ABRBKDQ.DATA,DISP=SHR
```

BACKUP ONLY SELECTED DATA SETS Do a manual backup (TYPE=DSF) of certain data sets using tape hardware compression (IDRC). The combination of ONLVOL and CATDSN causes only the volumes on which those data sets are cataloged to be processed. Using catalog data set name filtering ([See Section 80.14](#)), all ISPF data sets for "USERnn" and the current generation of any "GLEDGER" GDGs will be dumped. A listing of the selected data sets in ABR PRINT VTOC format will be printed.

Even though this is a TYPE=DSF backup, the backup files created will be part of the ABR backup subsystem. They will be assigned a cycle number in the current generation of the volumes involved, and will be read if an ABR full-volume restore is required. ABR Application Backup ([Section 52](#)) is a better technique for dumping selected data sets.

```
//DSF          EXEC   PGM=FDRABR,REGION=1124K
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=FDR,DISP=(,KEEP),UNIT=CART,DCB=TRTCH=COMP
//SYSIN       DD     *
              DUMP    TYPE=DSF,ONLVOL,PRINT=ABR
              SELECT  CATDSN=USER++.**ISPF**
              SELECT  CATDSN=GLEDGER.**(0)
```

50.20 CONTINUED . . .

**AUTOMATIC
INCREMENTAL
OR FULL
VOLUME
BACKUP**

Do an incremental backup of all updated data sets on all online disk volumes (same as DUMP TYPE=ABR), except that ABR will force a full-volume dump for any selected volume if the number of incremental backups in the current generation exceeds the limit specified for CYCLE in the ABR model DSCB. TAPE1 and TAPE2 cause two disks to be dumped concurrently, and TAPE11 and TAPE22 cause duplicate backups to be created. Incremental and full-volume backups will be mixed on the same set of tapes.

TYPE=AUTO can be used when you do not have a special time to do full-volume backups (such as weekends), but prefer to do full-volume dumps of certain volumes each night and incremental backups of the rest of the volumes. Once the cycle counts of the volumes have been synchronized so that only a selected subset of your volumes will exceed their count each night, the process is automatic. If you do prefer to do full-volume backups on weekends or another special time, use separate jobs with TYPE=FDR and TYPE=ABR instead.

```
//BACKUP      EXEC  PGM=FDRABR,REGION=3M
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//TAPE1       DD    UNIT=(TAPE,2),DSN=FDR1,DISP=(,KEEP),VOL=(,,255)
//TAPE11      DD    UNIT=(TAPE,2),DSN=FDR11,DISP=(,KEEP),VOL=(,,255)
//TAPE2       DD    UNIT=(TAPE,2),DSN=FDR2,DISP=(,KEEP),VOL=(,,255)
//TAPE22      DD    UNIT=(TAPE,2),DSN=FDR22,DISP=(,KEEP),VOL=(,,255)
//SYSIN       DD    *
              DUMP      TYPE=AUTO,ONLINE
```

**INCREMENTAL
BACKUP OF
SMS VOLUMES**

Dump any updated data sets from all volumes in two SMS storage groups, creating backups on one tape drive. Note that SMS volumes can also be selected by normal DISKxxxx DDs or VOL=/VOLG= operands. SMS management class attributes ([See Section 70](#)) may be used to exclude certain data sets from backup.

```
//BACKSMS     EXEC  PGM=FDRABR,REGION=2148K
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//TAPE1       DD    DSN=FDRSMS,UNIT=TAPE,DISP=(,KEEP),VOL=(,,99)
//SYSIN       DD    *
              DUMP      TYPE=ABR,SMSMANAGE=YES
              MOUNT     STORGRP=DBLARGE
              MOUNT     STORGRP=DBSMALL
```

**SIMULATE
INCREMENTAL
BACKUP**

Simulate the incremental backup of all online volumes beginning with "TSO". ABR will report on all updated data sets which would be selected by a real DUMP TYPE=ABR.

```
//SIM         EXEC  PGM=FDRABR
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//TAPE1       DD    DUMMY
//SYSIN       DD    *
              SIM       TYPE=ABR
              MOUNT     VOLG=TSO
```

50.20 CONTINUED . . .

**DUMP FROM
MULTIPLE
VOLUMES
CONCURREN-
TLY WITH
EXHPDM**

Backup many disk volumes concurrently using the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM is invoked by the SUBSYS= operands on the TAPE DD statements; [see Section 80.33](#) and the ExHPDM program documentation for details on the values to provide. Because there are 3 TAPE DD statements, ABR will backup 3 disk volumes concurrently but these concurrent backups will be interleaved into one file on a tape managed by ExHPDM. As each disk volume completes, ABR will select another volume to backup; these additional backups will be added to the same output tape file by ExHPDM, creating one large tape file containing all of the backups. This example illustrates full-volume backups (TYPE=FDR) but it could also be used for incremental backups (TYPE=ABR).

```
//DUMP          EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=BACKUP1,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN1      DD     SYSOUT=*
//TAPE2         DD     DSN=BACKUP2,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN2      DD     SYSOUT=*
//TAPE3         DD     DSN=BACKUP3,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRIN3      DD     SYSOUT=*
//SYSIN         DD     *
              DUMP    TYPE=FDR,MAXERR=1
              MOUNT   VOLG=DB
```


50.21 DATA SET RESTORE EXAMPLES

**RESTORE
SEVERAL
DATA SETS
FROM BACKUP**

Several data sets are to be restored from ABR Volume Backups; the most current backup copy of each is selected. The TAPE1 DD allocates a tape drive, using a dummy DSN and volume serial; ABR will use it to mount the correct backup tape for each restore (but it must be the right type of tape drive for each tape required).

- Data set A.B.C is cataloged on some volume; its backup from that volume will be restored, replacing the existing contents of the data set.
- Data set ACCT.PAYROLL is restored the same way, except that its backup is restored from the duplicate backup, COPY2.
- Data Set PAYROLL.MASTER is not cataloged, so VOL=MASTER specifies both the volume from which the backup was taken and the volume to which it will be restored.
- Data Set SYS1.XYZ is to be restored as SYS2.XYZ, which will be allocated and cataloged on one of the online volumes starting with SYS.

```
//RESTORE      EXEC   PGM=FDRABR,REGION=1024K
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DSN=FDR,VOL=SER=FDR,
//              UNIT=(TAPE,,DEFER),DISP=(OLD,KEEP)
//SYSIN         DD     *
RESTORE        TYPE=ABR
SELECT         DSN=A.B.C
SELECT         DSN=PAYROLL.MASTER,VOL=MASTER
S              DSN=ACCT.PAYROLL,COPY=2
S              DSN=SYS1.XYZ,NEWI=SYS2,NVOL=SYS*
```

**RESTORE
SEVERAL
DATA SETS
WITH
DYNTAPE**

Several data sets are to be restored from Volume Backups. They may be on a mixture of backup device types, so DYNTAPE is used to dynamically allocate the backup files as required. One data set will have its latest backup restored, the other will restore from a specific ABR generation and cycle of the volume it is currently cataloged to. In both cases, the data sets will be restored to the volume they are cataloged to, overlaying their existing contents.

```
//RESTORE      EXEC   PGM=FDRABR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
RESTORE        TYPE=ABR,DYNTAPE
SELECT         DSN=USER.FILE1
S              DSN=USER.FILE2,GEN=5,CYCLE=2
```

50.21 CONTINUED . . .

**PROCESS THE
ABR RESTORE
REMOTE
QUEUE**

Users have added requests for restore from Volume Backups to the ABR backup restore remote queue file. This JCL can be used to process those requests. The SYSIN file must contain a RESTORE TYPE=ABR statement. The ABRREST DD points to the remote queue file. ABR will ENQ on the remote queue, read the queued requests (if any), clear the queue, and immediately release it so that users are free to add new requests. Then it will process the requests, dynamically allocating the backup files as required. TSO users will receive messages on the success or failure of their request. If this queue is in use, this job should probably be run several times a day.

```
//REMOB    EXEC   PGM=FDRABR,REGION=1024K
//SYSPRINT DD     SYSOUT=*
//SYSIN    DD     *
           RESTORE TYPE=ABR,DYNTAPE
//ABRREST  DD     DSN=FDRABR.ABRREST.USER,DISP=SHR
```

**SIMULATE
DATA SET
RESTORE**

A restore of selected data sets from backup is simulated by ABR, identifying the backup volumes required. This might be used to pre-pull required tapes from the library.

```
//SIM      EXEC   PGM=FDRABR,REGION=1024K
//SYSPRINT DD     SYSOUT=*
//TAPE1    DD     DUMMY
//SYSIN    DD     *
           SIMREST TYPE=ABR
           SELECT  DSN=USERA.FILE1
           SELECT  DSN=USERA.FILE2
```

**RESTORE
FROM PRIOR
BACKUP**

A data set is to be restored, not from its most recent backup, but from an earlier backup. OLDBACKUP=2 says to restore from the third most recent backup (0 is the most recent). OLDBACKUP can be used only if the OLDBACKUP option is enabled for the volume from which this data set was dumped. The PRINT BACKUP statement of FDRABRP ([See Section 53](#)) can be used to display the backups of a data set and the dates they were taken.

```
//OLDBACK  EXEC   PGM=FDRABR,REGION=1024K
//SYSPRINT DD     SYSOUT=*
//SYSUDUMP DD     SYSOUT=*
//SYSIN    DD     *
           RESTORE TYPE=ABR,DYNTAPE
           SELECT  DSN=USER.DATA,OLDBACKUP=2
```

50.21 CONTINUED . . .

RESTORE MULTI-VOLUME DATA SETS VSAM and non-VSAM multi-volume data sets must be restored to the same number of volumes they were dumped from. There are no special procedures for restoring a multi-volume non-VSAM data set other than ensuring that ABR has sufficient volumes to allocate it on. When restoring to the data set's original volumes, this is automatic. When restoring to new volumes, the NVOL= parameter must specify enough volumes. The new volumes can also be provided by the ABR RESTORE ALLOCATION LIST. RECAT and VRECAT will insure that the restored data sets are cataloged to their new volumes. Data set USER33.MULTI.VOL will be restored as USER33.MULTIX.VOL

```
//MULTIVOL      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
  RESTORE       TYPE=ABR,DYNTAPE,RECAT,VRECAT
S              DSN=USER11.MULTI.VOL
S              DSN=USER22.MULTI.VOL,NVOL=(PROD01,PROD02,PROD03)
S              DSN=USER33.MULTI.VOL,NVOL=PROD*,NEWI=.MULTIX
```

See Section 80.13 for considerations for restoring multi-volume ICF VSAM clusters.

RESTORE DATA SETS MANUALLY In most cases, the records kept by ABR are sufficient to automatically restore the backup of a data set even when all you know is its data set name, as shown in the preceding examples. However, sometimes this information is lost (for example, if a volume is re-initialized with ICKDSF, all information about the data sets previously on it is lost). In this case, you will have to specify the volume that contained the data sets and the ABR generation and cycle of the backup containing the backup to be restored. Also, if you want to select data sets to be restored with a data set name mask, you must also specify VOL/GEN/CYCLE.

```
//RESTMAN       EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
  RESTORE       TYPE=ABR,DYNTAPE,RECAT,VRECAT
S              DSN=USER.DATA,VOL=TS0123,GEN=24,CYCLE=3
S              DSN=USER2**,VOL=TS0321,GEN=17,CYCLE=0
```

50.22 SMS DATA SET RESTORE EXAMPLES

There are no special considerations when SMS-managed data sets are restored in place, on top of their existing allocations; ABR just replaces their contents. When any data set being restored by ABR must be allocated on a system with SMS active, the SMS ACS routines will be invoked to decide if it is to be SMS-managed, to assign SMS classes, and to decide which volume it is to be placed on.

Any of the data sets in the examples in preceding [section 50.21](#) could have been SMS-managed. If they needed to be allocated, ABR would have passed their previous SMS classes (if any) to your ACS routines, which could accept, reject or override them. If an SMS storage class is assigned to a data set, SMS will select a storage group and a volume for the allocation. The following examples show how to influence or override the SMS allocation.

**OVERVERRIDE SMS
CLASSES**

Restore a data set from Volume Backups and request that the data set be SMS-managed; the original data set might be SMS-managed or not, but it is not currently cataloged; VOL= specifies the volume it was originally dumped from. The values specified for STORCLAS= and MGMTCLAS= will be passed to the SMS ACS storage and management class routines, which may override them. If your ACS routines accept your storage class or assigns a new storage class, SMS will then select a storage group and a volume. If your ACS routines assign a null storage group (non-SMS-managed) to the data set, it will be restored to the volume from which it was dumped (unless that volume was SMS-managed, since a non-SMS data set cannot be restored to a SMS volume).

```
//RESTSMS      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
RESTORE        TYPE=ABR,DYNTAPE
SELECT         DSN=PROD.MASTER.FILE,DATACLAS=MASTER,
               MGMTCLAS=PERM,STORCLAS=PROD2,VOL=PROD01
```

**RESTORE TO
NON-SMS
VOLUME**

Restore a data set from Volume Backups and request that the data set be non-SMS; the original data set might be SMS-managed or not. It does not exist on disk but is recorded in the ABR scratch catalog. NULLSTORCLAS specifies that a null value will be passed to the SMS ACS storage class routine, which may override it. If your ACS routines honor the null storage group, it will be restored as non-SMS on the volume TSO123; no SMS classes will be associated with it. If your ACS routines assign a storage class, SMS will then select a storage group and a volume.

```
//RESTSMS      EXEC   PGM=FDRABR,REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
RESTORE        TYPE=ABR,RECAT,DYNTAPE
SELECT         DSN=USER01.ISPF.ISPPROF,NULLSTORCLAS,NVOL=TSO123
/*
```

50.23 FULL VOLUME RESTORE EXAMPLES

Note that ABR full-volume restores can be enhanced if you are licensed for FDRCLONE, documented starting in [Section 50.70](#). FDRCLONE is particularly important if you are backing up to high-capacity tape cartridges, such as the IBM Magstar or StorageTek Redwood or 9840.

FDRCLONE allows you to restore only the data sets that are required at a disaster/recovery site or a test system, instead of restoring all data sets including those which will never be used on that system. FDRCLONE restores data sets from ABR backups on demand, as they are needed.

FDRDRP is part of FDRCLONE. It is a full-volume restore program, with the same result as an ABR full-volume restore, but it manages the backup tapes to minimize tape mounts and reduce the elapsed time of the restores. Normal ABR full-volume restore does one disk volume at a time, and may mount an input tape multiple times. FDRDRP does many disks in parallel and mounts input tapes a minimum number of times (often only once).

**FULL-VOLUME
RESTORE
FROM
CURRENT
BACKUP**

Two disk volumes are to be completely restored from ABR full-volume and incremental backups, beginning with the most recent incremental backup:

- PACK01 is to be restored back on top of itself. Its ABR model DSCB contains the most recent generation and cycle created.
- PACK02 has been destroyed, so its backup is to be restored onto a spare volume (SCR123). Since PACK02 is not online, ABR will search the ABR catalog for the most recent generation/cycle. SCR123 will be relabeled PACK02 at the end of the restore.

ABR will begin with the most recent incremental for each volume, and work back through the incrementals to the full-volume backup. The two volumes will be restored one at a time, in the order of the SELECT statements.

```
//RESTFULL      EXEC   PGM=FDRABR,REGION=1024K
//SYSPRINT       DD     SYSOUT=*
//SYSUDUMP       DD     SYSOUT=*
//DISK1          DD     UNIT=3390,VOL=SER=PACK01,DISP=OLD
//DISKX          DD     UNIT=3380,VOL=SER=SCR123,DISP=OLD
//SYSIN          DD     *
RESTORE          TYPE=FDR,CPYVOLID=YES,DYNTAPE
SELECT           VOL=PACK01
SELECT           VOL=PACK02,NVOL=SCR123
```

**RESTORE OF
FULL-VOLUME
BACKUP ONLY**

A disk volume is to be recreated using only the most recent full-volume backup recorded in the ABR catalog. The original volume SYSRES is being restored to a scratch volume SCR001. Since it will be relabeled back to SYSRES at a later time, VOLRESET=NO prevents ABR from renaming the ABR Model DSCB or the indexed VTOC.

This should be used when your installation has chosen to recover certain volumes using only the full-volume backup, rather than applying incremental backups. GEN=CURRENT forces ABR to locate the most recent generation in the catalog, while CYCLE=00 causes ABR to restore from only the full-volume backup which began that generation.

```
//RESTORE        EXEC   PGM=FDRABR,REGION=1M
//SYSPRINT       DD     SYSOUT=*
//SYSUDUMP       DD     SYSOUT=*
//SYSIN          DD     *
RESTORE          TYPE=FDR,VOLRESET=NO,DYNTAPE,ONLINE
SELECT           VOL=SYSRES,NVOL=SCR001,GEN=CURRENT,CYCLE=00
```

50.24 FDR INSTANTBACKUP EXAMPLES

BACKUP FROM SYMMETRIX BCV On an EMC Symmetrix subsystem, BCVs (Business Continuance Volumes) have been permanently assigned to the online payroll volumes PAYxxx; a previous one-time ESTABLISH has been issued to establish those pairings. Here are examples of full-volume, incremental, and selected data set backups.

Step FULL1 will split each BCV from its online volume to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

Step FULL2 (which can be run in a separate job) will do the actual backups. BCV=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the Symmetrix hardware to get the addresses of the BCVs; there is no need to specify them here.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY,LABEL=EXPDT=99000
//TAPE11      DD     DUMMY,LABEL=RETPD=60
//TAPE2       DD     DUMMY,LABEL=EXPDT=99000
//TAPE22      DD     DUMMY,LABEL=RETPD=60
//SYSIN       DD     *
      SPLIT TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
      MOUNT VOLG=PAY*
//FULL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN       DD     *
      DUMP TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,BCV=(USE,RET)
      MOUNT VOL=PAY*
```

50.24 CONTINUED . . .

BACKUP FROM SYMMETRIXBCV This example is the same as the previous Symmetrix BCV example, except that ABR incremental backups are taken:

(continued)

```
// INCR1      EXEC   PGM=FDRABR, REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY, LABEL=EXPDT=99000
//TAPE11      DD     DUMMY, LABEL=RETPD=60
//TAPE2       DD     DUMMY, LABEL=EXPDT=99000
//TAPE22      DD     DUMMY, LABEL=RETPD=60
//SYSIN       DD     *
      SPLIT TYPE=ABR, DSNENQ=USE, ENQERR=NO
      MOUNT VOLG=PAY*
// INCR2      EXEC   PGM=FDRABR, REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1, UNIT=CART, DISP=(MOD, KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1, UNIT=CART, DISP=(NEW, KEEP), LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2, UNIT=CART, DISP=(MOD, KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2, UNIT=CART, DISP=(NEW, KEEP), LABEL=RETPD=60
//SYSIN       DD     *
      DUMP TYPE=ABR, DSNENQ=USE, ENQERR=NO, BCV=(USE, RET)
      MOUNT VOL=PAY*
```

This example is the same as the previous Symmetrix BCV example, except that only selected data sets are backed up. This can be used to backup data sets belonging to one application with InstantBackup, but you cannot run more than one of these backups on a given disk volume at a time:

```
//SEL1       EXEC   PGM=FDRABR, REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY, LABEL=EXPDT=99000
//TAPE11      DD     DUMMY, LABEL=RETPD=60
//TAPE2       DD     DUMMY, LABEL=EXPDT=99000
//TAPE22      DD     DUMMY, LABEL=RETPD=60
//SYSIN       DD     *
      SPLIT TYPE=DSF, DSNENQ=USE, ENQERR=NO
      SELECT DSN=PAYROLL**
      MOUNT VOLG=PAY*
//SEL2       EXEC   PGM=FDRABR, REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1, UNIT=CART, DISP=(MOD, KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1, UNIT=CART, DISP=(NEW, KEEP), LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2, UNIT=CART, DISP=(MOD, KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2, UNIT=CART, DISP=(NEW, KEEP), LABEL=RETPD=60
//SYSIN       DD     *
      DUMP TYPE=DSF, DSNENQ=USE, ENQERR=NO, BCV=(USE, RET)
      SELECT DSN=PAYROLL**
      MOUNT VOL=PAY*
```

50.24 CONTINUED . . .

**BACKUP WITH
SNAPSHOT**

On a virtual disk subsystem (IBM RVA or StorageTek Iceberg or SVA) with the Snapshot hardware feature installed, snapshot target volumes have been permanently assigned to the online payroll volumes PAYxxx. Here are examples of full-volume, incremental, and selected data set backups.

Step FULL1 will snap each online volume to its snapshot target to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

Step FULL2 (which can be run in a separate job) will do the actual backups. SNAP=(USE,REL) tells ABR to determine if a snapped point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR remembers the addresses of the snapshot targets; there is no need to specify them here. DCT=YES requests that the HSDM (High Speed Data Mover) hardware option be used if it is installed on your disk subsystem; HSDM can reduce backup times up to 60%.

```
//FULL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY,LABEL=EXPDT=99000
//TAPE11      DD     DUMMY,LABEL=RETPD=60
//TAPE2       DD     DUMMY,LABEL=EXPDT=99000
//TAPE22      DD     DUMMY,LABEL=RETPD=60
//SYSIN       DD     *
              SNAP   TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
              MOUNT  VOL=PAY001,SNAPUNIT=07C8
              MOUNT  VOL=PAY002,SNAPUNIT=07C9
              MOUNT  VOL=PAY003,SNAPUNIT=07CA
//FULL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN       DD     *
              DUMP   TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,SNAP=(USE,REL),DCT=YES
              MOUNT  VOL=PAY*
```


50.24 CONTINUED . . .

**BACKUP WITH
SNAPSHOT
(continued)**

This example is the same as the previous Snapshot example, except that ABR incremental backups are taken:

```
// INCR1      EXEC  PGM=FDRABR , REGION=0M
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//TAPE1       DD    DUMMY , LABEL=EXPDT=99000
//TAPE11      DD    DUMMY , LABEL=RETPD=60
//TAPE2       DD    DUMMY , LABEL=EXPDT=99000
//TAPE22      DD    DUMMY , LABEL=RETPD=60
//SYSIN       DD    *
  SNAP      TYPE=ABR , DSNENQ=USE , DATA=USED , ENQERR=NO
  MOUNT     VOL=PAY001 , SNAPUNIT=07C8
  MOUNT     VOL=PAY002 , SNAPUNIT=07C9
  MOUNT     VOL=PAY003 , SNAPUNIT=07CA
// INCR2      EXEC  PGM=FDRABR , REGION=0M
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//TAPE1       DD    DSN=FDRABR . LASTAPE . ABR1 , UNIT=CART , DISP=(MOD , KEEP) ,
//              LABEL=EXPDT=99000
//TAPE11      DD    DSN=ABR1 , UNIT=CART , DISP=(NEW , KEEP) , LABEL=RETPD=60
//TAPE2       DD    DSN=FDRABR . LASTAPE . ABR2 , UNIT=CART , DISP=(MOD , KEEP) ,
//              LABEL=EXPDT=99000
//TAPE22      DD    DSN=ABR2 , UNIT=CART , DISP=(NEW , KEEP) , LABEL=RETPD=60
//SYSIN       DD    *
  DUMP      TYPE=ABR , DSNENQ=USE , ENQERR=NO , SNAP=(USE , REL)
  MOUNT     VOL=PAY*
```

50.24 CONTINUED . . .

**BACKUP WITH
SNAPSHOT
(continued)**

This example is the same as the previous Snapshot example, except that only selected data sets are backed up. This can be used to backup data sets belonging to one application with InstantBackup, but you cannot run more than one of these backups on a given disk volume at a time:

```
//SEL1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
  SNAP      TYPE=DSF,DSNENQ=USE,ENQERR=NO
  SELECT    DSN=PAYROLL**
  MOUNT     VOL=PAY001,SNAPUNIT=07C8
  MOUNT     VOL=PAY002,SNAPUNIT=07C9
  MOUNT     VOL=PAY003,SNAPUNIT=07CA
//SEL2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11     DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2      DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22     DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN      DD     *
  DUMP      TYPE=DSF,DSNENQ=USE,ENQERR=NO,SNAP=(USE,REL)
  SELECT    DSN=PAYROLL**
  MOUNT     VOL=PAY*
```

50.24 CONTINUED . . .

**BACKUP FROM
HDS
SHADOWIMAGE**

On an HDS 7700E subsystem, ShadowImage volumes have been permanently assigned to the online payroll volumes PAYxxx; a previous one-time PESTPAIR has been issued to establish those pairings. Here are examples of full-volume, incremental, and selected data set backups.

Step FULL1 will split each ShadowImage copy from its online volume to create point-in-time images and mark the backups as complete. It will create a new ABR backup generation and update the online volume with information about each new backup. As soon as this step completes, the point-in-time backup is complete and updates to the online volume can resume. Note that although the TAPE DD statements are DUMMY, they must specify the same EXPDT= or RETPD= values specified in the FULL2 step which follows.

Step FULL2 (which can be run in a separate job) will do the actual backups. PPRC=(USE,RET) tells ABR to determine if a split point-in-time image exists for each volume processed; if so, that image is backed up instead of the online volume. Volumes which have not been captured will be processed normally. ENQ=ON (the default) and DSNENQ=USE will be ignored when backing up from a point-in-time image. Note that ABR queries the ShadowImage hardware to get the addresses of the ShadowImage copies; there is no need to specify them here.

```
//FULL1      EXEC   PGM=FDRABR,REGION=OM
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY,LABEL=EXPDT=99000
//TAPE11      DD     DUMMY,LABEL=RETPD=60
//TAPE2       DD     DUMMY,LABEL=EXPDT=99000
//TAPE22      DD     DUMMY,LABEL=RETPD=60
//SYSIN       DD     *
      PSPLIT TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO
      MOUNT  VOLG=PAY*
//FULL2      EXEC   PGM=FDRABR,REGION=OM
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN       DD     *
      DUMP  TYPE=FDR,DSNENQ=USE,DATA=USED,ENQERR=NO,PPRC=(USE,RET)
      MOUNT VOL=PAY*
```

50.24 CONTINUED . . .

This example is the same as the previous HSD ShadowImage example, except that ABR incremental backups are taken:

```
// INCR1      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DUMMY,LABEL=EXPDT=99000
//TAPE11      DD     DUMMY,LABEL=RETPD=60
//TAPE2       DD     DUMMY,LABEL=EXPDT=99000
//TAPE22      DD     DUMMY,LABEL=RETPD=60
//SYSIN       DD     *
      PSPLIT TYPE=ABR,DSNENQ=USE,ENQERR=NO
      MOUNT VOLG=PAY*
// INCR2      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSPRIN2    DD     SYSOUT=*
//TAPE1       DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11      DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2       DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22      DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN       DD     *
      DUMP TYPE=ABR,DSNENQ=USE,ENQERR=NO,PPRC=(USE,RET)
      MOUNT VOL=PAY*
```

50.24 CONTINUED . . .

This example is the same as the previous HDS ShadowImage example, except that only selected data sets are backed up. This can be used to backup data sets belonging to one application with InstantBackup, but you cannot run more than one of these backups on a given disk volume at a time:

```
//SEL1      EXEC   PGM=FDRABR,REGION=OM
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DUMMY,LABEL=EXPDT=99000
//TAPE11     DD     DUMMY,LABEL=RETPD=60
//TAPE2      DD     DUMMY,LABEL=EXPDT=99000
//TAPE22     DD     DUMMY,LABEL=RETPD=60
//SYSIN      DD     *
      PSPLIT TYPE=DSF,DSNENQ=USE,ENQERR=NO
      SELECT DSN=PAYROLL**
      MOUNT VOLG=PAY*
//SEL2      EXEC   PGM=FDRABR,REGION=OM
//SYSPRINT   DD     SYSOUT=*
//SYSPRIN1   DD     SYSOUT=*
//SYSPRIN2   DD     SYSOUT=*
//TAPE1      DD     DSN=FDRABR.LASTAPE.ABR1,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE11     DD     DSN=ABR1,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//TAPE2      DD     DSN=FDRABR.LASTAPE.ABR2,UNIT=CART,DISP=(MOD,KEEP),
//              LABEL=EXPDT=99000
//TAPE22     DD     DSN=ABR2,UNIT=CART,DISP=(NEW,KEEP),LABEL=RETPD=60
//SYSIN      DD     *
      DUMP TYPE=DSF,DSNENQ=USE,ENQERR=NO,PPRC=(USE,RET)
      SELECT DSN=PAYROLL**
      MOUNT VOL=PAY*
```

50.30 VOLUME BACKUP ISPF SUPPORT

The ABR ISPF dialogs support data set restores from Volume Backups. Users can use them to submit their own ABR jobs to restore specific data sets from the Volume Backups, or they can add the restore requests to a Remote Queue data set (described in [Section 50.03](#)).

The ISPF dialogs also support adding backup requests to a Remote Queue data set.

**ABR ISPF MAIN
MENU**

Innovation recommends that the ISPF dialogs be installed so that a character of "A" on your regular ISPF main menu will call the ABR main menu. Details of this installation are found in [Section 90](#). When this manual refers to a panel name such as A.S, it assumes that this convention was used. If your installation used a different letter on your ISPF main menu, or uses some other technique to get to the ABR main menu, you will have to change our examples appropriately.

If you enter simply "A" (or whatever) from your ISPF main menu, the ABR main menu is displayed:

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- FDR PRIMARY OPTIONS MENU -----
OPTION  ===>

    1  REPORTS      - ABR REPORTING FUNCTIONS
    2  RESTORE      - ABR DATA SET RESTORE
    3  ARCHIVE      - ABR DATA SET ARCHIVE OR SUPERSCRATCH
    4  BACKUP       - ABR DATA SET BACKUP
    5  REMOTE Q     - ABR REMOTE QUEUE UTILITY FUNCTIONS

    C  COMPAKTOR    - COMPAKTOR MAP AND SIMULATION REPORTS
    R  RELEASE      - COMPAKTOR RELEASE

    I  INSTALL      - INSTALLATION AND MAINTENANCE OF FDR AND OPTIONAL PRODUCTS
    J  JCL PARMS    - SPECIFY FDR JCL AND SYSOUT DEFAULTS FOR SUBMITTED JOBS
    K  FORMAT       - MODIFY FORMAT OF GENERATED REPORTS

    M  MESSAGES     - FDR MESSAGES AND CODES QUERY FACILITY
    Q  QUERY        - FDR/ABR STATISTICS QUERY

    S  SRS          - SEARCH, REPORT, SERVICES DIALOG

    T  FORTSEL      - BACKUP FILE MANAGEMENT UTILITY

```

Note that option 2 (RESTORE) can be used to restore the most recent backup of individual data sets from Volume Backups. Depending on the customization of the dialogs done during install ([Section 90](#)), the user may have the option to directly submit an ABR restore job or to add the request to a remote queue. Although the RESTORE panels supports restore of older backups, the SRS dialog, described next, is a far superior way of restoring data sets.

50.30 CONTINUED . . .

**SRS ISPF
DIALOG**

SRS stands for "Search, Report, and Services". It is described in detail in [Section 54](#). It allows users to search for information on data sets from a variety of sources, display the information in a easy-to-use format, and optionally to execute various TSO and ABR functions against the data sets displayed.

For Volume Backups, it can display backup information for the data sets selected. If OLDBACKUP was enabled on the volume containing the data sets, it can display all of the recorded backups for each data set, including the date of the backup. Simply by typing the RESTORE command on the proper line, users can request that a data set be restored from a particular backup.

Users can also add a request to the backup Remote Queue, so that the displayed data set will be included in the next incremental backup of the volume it resides on.

From your ISPF main menu, if you enter "A.S.1", you will receive this SRS panel:

In this example, the user has filled in several data set name masks, has specified several attributes of the selected data sets to display (including ABR backup information), and had requested that all recorded backups of each data set be displayed (OLDBACKUP=ALL).

```

----- DATASET SELECTION: ABRBKUP ----- LINE 1/16 COL 4:6/10
COMMAND ===>                                SCROLL ===> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD          SELECTION VALUE                                REPORT SORT
-----
DSNAME  ===>  bab.*.jcl  bab.ac**_----- 1_-- --
VOL      ===>  ----- 2_-- --
SOURCE   ===>  CATALOG (Catalog Volume Archive Appl Scratch Extract)  --- --
CATALOGN ===>  -----
DEVTYPE  ===>  ----- 3_-- --
DSORG    ===>  ----- 4_-- --
BKGEN    ===>  ----- 6_-- --
BKCYCLE  ===>  ----- 7_-- --
BKINFO   ===>  ----- 8_-- --
OLDBKUP  ===>  ALL_----- --
SIZE     ===>  ----- 5_-- --

```

Note that all of the SRS report and selection fields shown may not display on the screen at once, you may need to scroll up and down to find the SRS fields you need. [Section 54](#) includes information and examples on customizing SRS to display just the information required for a particular function, such as restore from Volume Backup.

50.30 CONTINUED . . .

In response to this request, SRS will search the system catalogs for the data sets requested and display their backup information in a format similar to:

----- DATASET LIST: DSLIST ----- LINE 1/20 COL 3:9/12									
COMMAND ==>					SCROLL ==> HALF				
3 DATA SETS SELECTED.									
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help
COMMAND	DATA SET NAME				VOLSER	DEVTYPE	ALLOC	DSO	BKGN BKC BKDATE

	BAB.HCHECKV2.CNTL				IDPLB0	ARCHIVED			
	BAB.JCL.CNTL				TSOWK0	3380	150	PO	230 1 96.162
									230 0 96.159
									229 4 96.158
									229 2 96.156
									229 1 96.155
									229 0 96.152
restore									228 3 96.151
									228 2 96.150
									228 1 96.149
									228 0 96.145
									227 3 96.143
									227 2 96.142
									227 0 96.138
									226 4 96.137
restore	BAB.AC.DATA				IDPLB3	3380-K	1	PS	721 0 96.159

For data set BAB.AC.DATA, the most recent 13 backups are displayed, with the generation (BKGN), cycle (BKC) and backup date (BKDATE). To restore from a particular backup, simply type RESTORE on the line next to the backup you want. For data set BAB.AC.DATA only one backup is shown; type RESTORE next to it to restore the most recent backup.

The RESTORE command will take you to another panel where you can set options for the restore:

```
----- FDRSRS - Backup Restore ----- Row 1 to 2 of 2
COMMAND ==> edit                                SCROLL ==> HALF

Edit JCL    Submit JCL    FG - execute in the foreground    RQ - add to remote q

Operands for RESTORE TYPE=ABR statement (section 50):
==> RESTORE TYPE=ABR,DT

DSNAME/Filter ==> 'BAB.JCL.CNTL'                                Gen ==> 0228
Volume Serial ==> TSOWK0                                        Cycle => 03
New DSNAME    ==>                                              OLDB  =>
or NEWINDEX   ==>                                              Copy  => 1
New Volser(s) ==>

Operands for SELECT DSN= statement (section 50):
==> NOTIFY=BAB

DSNAME/Filter ==> 'BAB.AC.DATA'                                Gen ==> 0721
Volume Serial ==> IDPLB3                                        Cycle => 00
New DSNAME    ==>                                              OLDB  =>
or NEWINDEX   ==>                                              Copy  => 1
New Volser(s) ==>

Operands for SELECT DSN= statement (section 50):
==> NOTIFY=BAB
```


50.30 CONTINUED . . .

The proper volume, generation, and cycle have already been filled in. If required, you can enter a new name or NEWINDEX for renaming the restored data set; you can also specify a new target volser if it is not to be restored to the original volume.

The user has the option to submit a batch ABR jobstream to perform the restore (EDIT or SUBMIT commands), to attach ABR under TSO and execute the restore in the foreground (FG command), or to add the requests to the ABR Restore Remote Queue (RQ command, only valid if you have implemented that remote queue as described in [Section 90](#)). Note that foreground (FG) restores will work only if the TSO user is authorized to mount tapes. Your installation may have modified this panel to limit the choices available, or it may execute a particular command (e.g., RQ) automatically without displaying this panel.

Because of the NOTIFY= operand, your TSO session will receive a message when the restore of each data set is complete.

50.40 VTOC MAINTENANCE UTILITY (FDRABRM)

FDRABRM is an ABR utility program which makes modifications in the VTOCs of DASD volumes to enable ABR processing and request special ABR processing.

WHY A VTOC UTILITY

Before ABR processing of a DASD volume (other than Application Backup) can occur, a special DSCB (data set) must exist in the VTOC of that volume. This is known as the ABR Model DSCB; it is called a model since the data set has no space (0 tracks) assigned, just like the "model DSCBs" used with GDGs. ABR uses fields in the ABR Model DSCB to store processing options for the volume, and to record information about recent full-volume and incremental backups. FDRABRM can create and modify the ABR Model DSCB. The data set name of the ABR Model DSCB is "FDRABR.Vvolser" (the prefix FDRABR may be changed by your installation at the time that ABR is installed).

Information about recent full-volume and incremental backups of individual data sets is stored in the Format 1 DSCBs of those data sets. It is also possible to set special ABR processing options for specific data sets in the DSCB. FDRABRM initializes and sets the DSCB fields used by ABR.

FDRABRM STATEMENTS

ABRINIT - enables a DASD volume for ABR processing by creating the ABR Model DSCB and setting the initial processing options for the volume. ABRINIT can be executed either on a new volume that has just been initialized by ICKDSF, or on an existing volume that contains data sets. **"ABR initialization" does not affect any data already on the volume.**

MAINT - changes the ABR processing options in the ABR Model DSCB.

REMODEL - recreates the ABR Model DSCB when it is accidentally destroyed or improperly modified and resynchronizes the generation and cycle numbers.

DEFAULT - sets defaults for following FDRABRM statements. It is not documented in this manual, but it accepts many operands from the other statements, including CYCLE=, GEN=, HIGHTHRESOLD=, LIST=, LOWTHREASHOLD=, RESERVE=, RETPD=, and RETPD2=. DEFAULT can be used when you have several following statements that will require the same operands; you can enter them once on DEFAULT and remove them from the individual statements.

NOTE: Most of the functions of the ABRINIT, MAINT, and REMODEL statements may be conveniently executed using the ABR ISPF dialog (option A.I.8). This is the recommended method for initializing and maintaining ABR volume processing options. See Section 50.46 for details.

DSCB FIELDS

In the Format 1 DSCB of each data set, ABR uses:

- the Update Indicator - a flag set by data management whenever the data set is opened for OUTPUT or UPDATE and reset by ABR when a data set is backed up by full-volume or incremental dumps.
- the Last Reference Date - set by data management whenever the data set is opened, if it is already non-zero. Since the Last Reference Date is normally set to zero when a data set is allocated, FDRABRM has an option to set zero LRDs to today's date, to initiate updating of that field. However, the ABR DADSM Pre-processing exit ([Section 90](#)) also sets the Last Reference Date to the creation date when the data set is created, so the FDRABRM function is usually not required unless the DADSM exit is not enabled.
- two additional bytes (relative locations 103 and 104) to store information about the ABR status of the data set.
- If the OLDBACKUP option is enabled on the volume, an additional field to store information on pervious backups.

50.40 CONTINUED . . .

**ABR
PROCESSING
OPTIONS**

FDRABRM may be used to set these ABR options for an entire DASD volume:

- Enable or disable ARCHIVING (DUMP TYPE=ARC)
- Enable or disable SUPERSCRATCH (DUMP TYPE=SCR)
- Enable or disable recording of the location of backups other than the most recent (OLDBACKUP)
- Volume allocation thresholds for ARCHIVE processing

**UNDOCU-
MENTED
FUNCTIONS**

FDRABRM has several undocumented functions. They are rarely used and have been eliminated from the manual, although they are still accepted. Contact Innovation if you have a need for these functions.

FDRABRM is capable of setting ABR options for individual data sets, such as "always backup", "never backup", and "never ARCHIVE". Since the data set options are stored in the DSCB of the data set involved, the options will be lost if the dataset is deleted and reallocated. Innovation recommends that you use other means to accomplish the same results, such as the ABR protect lists and SELECT statements in the ABR input. FDRABRM can also be used to change the "update" flag and last reference date in individual DSCBs, but this function is rarely needed. For these reasons, the SELECT statement for setting these options is no longer documented in the manual, although it is still accepted.

FDRABRM has one additional maintenance function: information about old backups (prior to the most recent) is stored in one of several reserved fields in the Format 1 DSCB; the choice is controlled by the user in the FDR Global Option table, although there is currently only one choice (DS1SYSCD, the System Code field) that is valid on all systems. If your installation has previously used a different field for old backup information, contact Innovation for options which allow you to change to use of DS1SYSCD without losing any old backup information.

50.40 CONTINUED . . .

**SMS
CONSIDER-
ATIONS**

The ABRINIT and REMODEL statements support the creation of ABR Model DSCBs on volumes managed by SMS (System Managed Storage; [see Section 70](#)). The following rules and considerations apply:

- If any volume to be processed by ABR is SMS-managed, then the ABR catalog must be an ICF catalog, since all data sets on SMS-managed volumes must be cataloged in an ICF catalog, including the ABR Model DSCB.
- Since the whole point of allocating an ABR Model DSCB is to put it on a specific DASD volume, the SMS Storage Group ACS routine will not be called to select a volume. The volumes on which ABR Model DSCBs are created will be selected based on the VOL, VOLG, or ONLINE operands or DISKxxxx DD statements.
- When creating an ABR Model DSCB on an SMS-managed volume a storage class name must be assigned. Your installation might want to set up its storage class ACS routine to assign an appropriate storage class to ABR Model DSCBs (data sets named FDRABR.Vxxxxxx with zero space); in this case, no special operands are needed on the ABRINIT or REMODEL statement. For an example, see "ABR Volume Selection" in [Section 70.12](#). If not, then a storage class must be assigned by the STORCLAS= operand of the ABRINIT or REMODEL statement; your installation's ACS routine may accept the requested storage class, or may change it to a different storage class. In these cases the ABR Model DSCB will be successfully created. It does not matter to FDRABRM what storage class is assigned, only that it is a valid storage class name.
- Your installation's ACS routine may change the requested storage class to null (blanks) which indicates that this data set should not be SMS-managed. If the target volume is SMS-managed, a storage class is required, so you must specify BYPASSACS on the ABRINIT or REMODEL statement to allow the ABR Model DSCB to be created with the storage class you specified (STORCLAS=).
- The ABR Model DSCB on an SMS-managed volume may optionally have a data class assigned. The data class may be assigned only by the DATACLAS= operand; the data class ACS routine is not given control. The data class has no effect on the processing of the ABR Model DSCB; if specified, it is used only for documentation or reporting.
- The ABR Model DSCB on an SMS-managed volume may optionally have a management class assigned. The management class may be assigned automatically by the ACS routine, or explicitly by the MGMTCLAS operand (unless overridden by the ACS routine). The specifications in the SMS definition for the management class will not be used by ABR for managing the ABR Model DSCB; the ABR Model DSCB will be backed up on every incremental backup, and will never be ARCHIVED. However, if your installation uses DASD management software that looks at the management class, the management class may be significant and should prevent the ABR Model from ever being migrated or deleted.

50.41 FDRABRM JCL REQUIREMENTS

The following Job Control Statements are required to execute the ABR VTOC maintenance utility:

EXEC STATEMENT Must specify the name of the ABR VTOC Maintenance Utility program, FDRABRM. The REGION=, if specified, must be at least 256K. The first or only control statement may be specified as PARM= data.

STEPLIB OR JOBLIB DD STATEMENT If required, must specify the load module library in which ABR resides. This library must be APF authorized.

SYSPRINT DD STATEMENT Specifies the output message data set. This is a required DD statement and usually is a SYSOUT data set.

ABRMAP DD STATEMENT Specifies the destination for the VTOC/DSCB map data set, produced if the LIST= operand is specified.. This is an optional DD statement and usually is a SYSOUT data set. If not present, the VTOC/DSCB map will be routed to SYSPRINT.

SYSUDUMP DD STATEMENT Specifies the abend dump data set. Although not required, we strongly urge you to always include this DD statement so that we can help you diagnose your error conditions. Usually a SYSOUT data set.

If you have the ABEND-AID product from COMPUWARE include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

SYSIN DD STATEMENT Specifies the control statement data set, usually an input stream or DD * data set. It can be omitted if the only control statement is specified by PARM= on the EXEC statement.

50.42 FDRAARM ABRINIT STATEMENT

ABRINIT ABR	BYPASSACS	,LIST=ALL YES <u>NO</u>
	,CYCLE=nn	,LOWTHRESHOLD=nnn
	,DATACLAS=classname	,MGMTCLAS=classname
	,DISABLE=(option,option,..)	,ONLINE
	,ENABLE=(option,option,..)	,RESERVE=NO <u>YES</u>
	,FORCE	,RETPD=nnnn
	,GEN=nnnn	,RETPD2=nnnn
	,HIGHTHRESHOLD=nnn	,STORCLAS=classname
		,VOL=vvvvvv
		,VOLG=vvvvvv

ABRINIT STATEMENT The ABRINIT statement is used to perform ABR initialization on a volume which is **not** currently initialized for ABR processing, one that has no ABR Model DSCB and which has **never been backed up** by ABR. It specifies the ABR disk volume processing options and the volumes for which they are to be set. If the ABRINIT statement contains the VOL or VOLG operand, the options will be set on only the specified volume(s). If the ABRINIT statement contains the ONLINE operand, the options will be set on all mounted volumes. In either case, only volumes not currently initialized for ABR will be altered.

NOTE: the presence of an ABR Model DSCB automatically enables the volume for ABR Volume Backups, as described in [Section 50.01](#). ENABLE= and DISABLE= set options in the ABR Model DSCB which enable or disable Archive Backups (DUMP TYPE=ARC) and Superscratch (DUMP TYPE=SCR), as described in [Section 51](#). ABR initialization is not required for Application Backup (DUMP TYPE=APPL, [see Section 52](#)).

OPERANDS BYPASSACS On a system with SMS (System Managed Storage) active, specifies that the SMS ACS (Automatic Class Selection) routines are not to be invoked when allocating the ABR Model DSCB. The SMS classes specified by the DATACLAS=, MGMTCLAS=, and STORCLAS= operands will be assigned if the volume on which the ABR Model DSCB is being created is SMS-managed. BYPASSACS must be specified if the installation's ACS routines do not permit a storage class to be assigned to an ABR Model DSCB on an SMS-managed volume.

Since BYPASSACS bypasses normal SMS allocation controls and rules, the user running the ABRINIT job must be authorized to the RACF profile STGADMIN.ADR.RESTORE.BYPASSACS in class FACILITY, or the equivalent in other security systems.

The default is that the SMS ACS routines will be invoked for the ABR Model DSCB. The SMS classes, if any, specified by the MGMTCLAS and STORCLAS operands will be passed to the ACS routines, which may accept, change, or nullify them.

50.42 CONTINUED . . .

CYCLE=	<p>Specifies the number of cycles that are to be created in each generation of backups for this volume at installations that execute program FDRABR with DUMP Statements that specify TYPE=AUTO. (For a full explanation of Cycles and Generations, See Section 50.01). The number specified must be between 0 and 63, inclusive.</p> <p>The default is 10 unless overridden by a preceding DEFAULT statement or by a change in the FDR Global Option Table (ISPF option A.I.4.7).</p>
DATACLAS=	<p>Specifies the SMS data class to be associated with ABR Model DSCBs that are created on SMS-managed volumes by this ABRINIT statement. The data class has no effect on the processing of the ABR Model DSCB; if specified, it is used only for documentation or reporting.</p> <p>The default is that no data class will be associated with the ABR Model DSCBs.</p>
DISABLE=	<p>The option(s) specified will be disabled for ABR processing on this volume. See the ENABLE Operand for the available options and their defaults.</p>
ENABLE=	<p>The option(s) specified will be enabled for ABR processing. Multiple options are enclosed in parentheses and separated by commas. The options are:</p> <p>ARCHIVE sets the indicator on this volume to permit data set ARCHIVING. If the program FDRABR is executed against this volume with a DUMP TYPE=ARC statement, the volume will be processed.</p> <p>Default is that the volume is enabled for ARCHIVE.</p> <p>OLDBKUP sets the indicator on this volume to permit tabling of previous backups, up to 13. If the program FDRABR is executed against this volume with a DUMP TYPE=ABR, TYPE=DSF, TYPE=FDR or TYPE=AUTO statement, the previous backup cycles will be recorded, in addition to the most recent backup. The location within the Format 1 DSCB of the old backup table as well as the maximum number of cycle table entries is specified stored in the FDR Global Option Table (See Section 90).</p> <p>Default is that the OLDBACKUP information will not be recorded.</p> <p>SCRATCH sets the indicator on this volume to permit SUPERSCRATCH (DUMP TYPE=SCR).</p> <p>Default is that the volume is disabled for SUPERSCRATCH.</p>

50.42 CONTINUED . . .

FORCE

FDRABRM verifies during initialization that the reserved fields (bytes 103 and 104) which will be used by ABR are currently binary zero in every Format 1 DSCB. If this volume has previously been processed by ABR but the ABR Model DSCB was accidentally destroyed, this condition may exist but you should use the REMODEL statement instead of ABRINIT to rebuild the ABR Model DSCB.

FORCE indicates that ABR initialization is to proceed even if one or more DSCBs have non-zero values in those bytes (they will be set to zero). Be cautious with the use of **FORCE**; do not specify it unless an ABRINIT without **FORCE** has failed because of non-zero reserved bytes, and you are sure that the information in those bytes is not needed.

The default is the initialization will fail if any DSCB has non-zero data in bytes 103 or 104.

Note: **FORCE** can and must be used if you have moved ABR-managed data sets to this volume from other volumes, since this may cause the ABR indicators to be set.

GEN=

Specifies the number of generations of backups for this volume that ABR is to maintain in the ABR catalog. At the beginning of a new generation, program FDRABR will uncatalog all the backup tapes from the nnnn'th prior generation, where nnnn was the value specified for **GEN=**. (For a full explanation of generations, [See Section 50.01](#)). The number specified must be between 1 and 1000, inclusive.

The default is 4 unless overridden by a preceding **DEFAULT** statement or by a change in the FDR Global Option Table (ISPF option A.I.4.7).

HIGHTHRESHOLD=

Specifies the high allocation percentage threshold that ABR is to store in the ABR Model DSCB. It may optionally be used to bypass **ARCHIVE** and **SUPERSCRATCH** on volumes which are below this threshold (For an explanation of threshold usage, [See Section 51](#)). The number specified must be between 0 and 100, inclusive.

The default is 80 unless overridden by a preceding **DEFAULT** statement or by a change in the FDR Global Option Table (ISPF option A.I.4.7).

LIST=

YES requests that the ABRINIT statement print a VTOC listing showing all data sets on the volumes being initialized.

ALL is the same as **YES** for ABRINIT.

NO specifies that the ABRINIT statement will not print a VTOC listing.

The default is **NO** unless overridden by a preceding **DEFAULT** statement.

LOWTHRESHOLD=

Specifies the low allocation percentage threshold that ABR is to store in the ABR Model DSCB. It may optionally be used to bypass **ARCHIVE** and **SUPERSCRATCH** on volumes which are below this threshold (For an explanation of threshold usage, [See Section 51](#)). The number specified must be between 0 and 100, inclusive.

The default is 50 unless overridden by a preceding **DEFAULT** statement or by a change in the FDR Global Option Table (ISPF option A.I.4.7).

50.42 CONTINUED . . .

MGMTCLAS=	<p>Specifies the SMS management class to be associated with ABR Model DSCBs that are created on SMS-managed volumes by this ABRINIT statement. The ACS routine may accept, change, or nullify this class, unless BYPASSACS is specified. The specifications in the SMS definition for the management class will not be used by ABR for managing the ABR Model DSCB; the ABR Model DSCB will be backed up on every incremental backup, and will never be ARCHIVED. However, if other DASD management software that looks at the management class runs on this volume, the management class may be significant; you should assign a class that prevents the ABR Model DSCB from ever being migrated or deleted.</p> <p>The default is that no management class will be associated with the ABR Model DSCBs, unless one is assigned by the ACS routine.</p>
ONLINE	<p>Specifies that this ABRINIT statement is to initialize all DASD volumes that are online to this system, except those volumes referenced by another ABRINIT statement in this run containing VOL= or VOLG= and those that are already initialized for ABR.</p> <p>If ONLINE is specified on a statement that also specifies VOL= or VOLG=, then ONLINE is ignored.</p>
RESERVE=	<p>YES specifies that a RESERVE statement is to be used to insure VTOC integrity during the update operations, in addition to an ENQ on SYSVTOC for the volume involved.</p> <p>NO indicates that only the SYSVTOC ENQ is to be issued.</p> <p>Default is YES unless specified on a preceding DEFAULT statement.</p>
RETPD= RETPD2=	<p>Specifies the default retention period, in days, for the full-volume backups that begin each new generation of backups for this volume. RETPD= is the default for TAPEX DDs (COPY 1) and RETPD2= for TAPEXX DDs. All cycles (incremental backups) within a generation will by default be set to expire on the same day as the preceding full-volume backup. Retention periods can be overridden at the time of the backups. For a full explanation of cycles, generations, and retention periods for backups, See Section 50. The number specified for RETPD or RETPD2 must be between 1 and 9999, inclusive.</p> <p>The default is 60 unless overridden by a preceding DEFAULT statement or by a change in the FDR Global Option Table (ISPF option A.I.4.7).</p>
STORCLAS=	<p>Specifies the SMS storage class to be associated with ABR Model DSCBs that are created on SMS-managed volumes by this ABRINIT statement. The ACS routine may accept, change, or nullify this class, unless BYPASSACS is specified. However, if no storage class is assigned for an ABR Model DSCB to be created on an SMS-managed volume, the creation will fail. See "SMS Considerations" in Section 50.40.</p> <p>The default is that no storage class will be associated with the ABR Model DSCBs, unless one is assigned by the ACS routine.</p>

50.42 CONTINUED . . .

VOL=

Specifies the serial number of the disk volume to be initialized by this ABRINIT statement. VOL= will accept only a single volume serial; it will not accept a list enclosed in parentheses, or a volume group indicated by placing an asterisk at the end of a prefix. For example, VOL=TEST01

VOLG=

Specifies a group of volume serial numbers. This ABRINIT statement will initialize any disk volume that is online to the system and that has a serial number starting with the 1 to 5 characters specified. For example, VOLG=TEST

NOTE: VOL= and VOLG= are mutually exclusive. One and only one of these operands may be specified.

50.43 FDRABRM MAINT STATEMENT

MAINT		,LIST=ALL YES NO
	CYCLE=nn	,LOWTHRESHOLD=nnn
	,DISABLE=(option,option,..)	,ONLINE
	,ENABLE=(option,option,..)	,RESERVE=NO YES
	,FORCE	,RESET=AUTOCYCL
	,GEN=nnnn	,RETPD=nnnn
	,HIGHTHRESHOLD=nnn	,RETPD2=nnnn
		,VOL=vvvvv
		,VOLG=vvvvv

MAINT STATEMENT The MAINT statement is used to change the ABR processing options currently set for a volume already initialized for ABR processing. For operands which specify ABR options, the defaults are the values currently stored in the ABR model DSCBs on the volumes involved. The MAINT statement **must** contain either the VOL= or VOLG= operand.

NOTE: If you change the volume serial number of a disk through any method other than a full-volume restore with ABR or FDR, you must run a MAINT statement before taking any ABR backups. The MAINT statement will adjust the data set name of the ABR model DSCB to reflect the new serial number of the disk. No operands except VOL= or VOLG= are required for this function.

OPERANDS All operands of MAINT, with the except of RESET=, described below, are also operands of the ABRINIT statement. Please see [Section 50.42](#) for descriptions of these operands.

RESET= **AUTOCYCL** sets the contents of the auto-cycle number to the current (latest) cycle number ([see Section 50.01](#)).

50.44 FDRABRM REMODEL STATEMENT

REMODEL	AUTOCYCLE=nn	,HIGHTHRESHOLD=nnn
	,BYPASSACS	,LASTGEN=nnnn
	,CYCLE=nn	,LASTCYCLE=nn
	,DATACLAS=classname	,LIST=ALL NO YES
	,DISABLE=(option1,option2,...)	,LOWTHRESHOLD=nnn
	,ENABLE=(option1,option2,...)	,MGMTCLAS=classname
	,EXPD=yyddd yyyyddd	,RETPD=nnnn
	,EXPD2=yyddd yyyyddd	,RETPD2=nnnn
	,FORCE	,STORCLAS=classname
	,GEN=nnnn	,VOL=vvvvvv

REMODEL STATEMENT The REMODEL statement recreates the ABR Model DSCB after it has been accidentally scratched or modified on a volume that was previously initialized for ABR. It can also be used to reset ABR volume backup information if the ABR Model DSCB is backleveled (restored by other than FDR or ABR). A separate REMODEL statement is required for each volume to be processed.

REMODEL accepts most of the operands that ABRINIT accepts, plus some additional operands used to specify current ABR backup information. REMODEL rebuilds the ABR Model DSCB, but does not do anything to the other format 1 DSCBs in the VTOC.

NOTE: If you restore a volume with SAR (Stand-Alone Restore, [see Section 15](#)), the ABR backup information in the ABR Model DSCB will be backleveled. The REMODEL statement should be used to reset the ABR Model DSCB to the current generation and cycle numbers indicated in the catalog. No operands except VOL and FORCE are required for this function.

OPERANDS All operands of REMODEL, with the except of AUTOCYCLE=, EXPD=, EXDP2=, LASTGEN= and LASTCYCLE=, described below, are also operands of the ABRINIT statement. Please [see Section 50.42](#) for descriptions of these operands. The FORCE operand has a slightly different meaning here, so it is also described below.

AUTOCYCLE= Specifies the most recent autocycle number for this volume. The number specified must be in the range of 0 to 63, inclusive.

If not specified, ABR will search the ABR catalog for the most current cycle and set AUTOCYCLE to that value.

50.44 CONTINUED . . .

EXPD= EXPD2=	Specifies the expiration date of the latest full volume backup (TYPE=FDR) as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are assumed to be in the 21st century, e.g., 12015=2012.015. For readability, a period may be inserted between the year and day (e.g., 1997.123). EXPD= is the expiration of the latest COPY1 backup, and EXPD2= is for COPY2. If either or both is not specified, ABR will set the expiration by adding the retention period to today's date. These expiration dates are used by ABR when creating incremental backups if no expiration or retention period is specified for them.
FORCE	Specifies that ABR is to update the Model DSCB if it currently exists on the disk volume. If not specified, ABR will fail if the volume already contains an ABR Model DSCB.
VOL=	Specifies the serial number of the disk volume on which to recreate the ABR Model DSCB. The VOL= operand is required.

50.45 VTOC UTILITY JCL EXAMPLES

NOTE: Most of the functions of the ABRINIT, MAINT, and REMODEL statements may be conveniently executed using the ABR ISPF dialog as shown in the next section.

**INITIALIZE ALL
ONLINE
VOLUMES I**

Set the ABR disk volume processing options for all volumes that are online to the system, enabling all ABR features including SUPERSCRATCH and OLDBACKUP. Use the defaults of CYCLE=10, RETPD=60, GEN=4, LOWTHRESHOLD=50, HIGHTHRESHOLD=80 unless these have been changed in the FDR Global Option Table. Volumes that are already initialized for ABR will be bypassed, as will SMS-managed volumes since no storage class is specified. Do not print the VTOCs.

```
//INITVOLS EXEC PGM=FDRABRM,REGION=256K
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
        ABRINIT          ONLINE,ENABLE=(OLDBKUP,SCRATCH)
```

**INITIALIZE ALL
ONLINE
VOLUMES II**

Set the ABR disk volume processing options for volume SYSRES with GEN=12, RETPD=90 and with ARCHIVE disabled. Set the disk volume processing options for all other volumes that are online to the system with GEN=4, RETPD=30 and ARCHIVE enabled. All volumes will have OLDBKUP enabled. Low and high allocation percentage thresholds of 40 and 75 will be set on all volumes. If the reserved DSCB bytes used by ABR are not binary zeroes on any volume (except SYSRES), FORCE cause the initialization to proceed, zeroing those bytes. Print the VTOCs of every volume (LIST=ALL).

```
//INIT EXEC PGM=FDRABRM,REGION=256K
//SYSPRINT DD SYSOUT=*
//ABRMAMP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
        DEFAULT LIST=ALL,LOWTHRESHOLD=40,HIGHTHRESHOLD=75
        ABRINIT VOL=SYSRES,GEN=12,RETPD=90,DISABLE=ARCHIVE,ENABLE=OLDBKUP
        ABRINIT ONLINE,RETPD=30,ENABLE=OLDBKUP,FORCE
```

**INITIALIZE
GROUPS OF
VOLUMES**

Set the ABR disk volume processing options for all of the TEST volumes with SUPERSCRATCH (TYPE=SCR) enabled. Set the ABR disk volume processing options for all of the PROD volumes with SUPERSCRATCH disabled. All volumes should have OLDBKUP enabled. 3 generations of backups will be kept for the TEST volumes, and 6 for the PROD volumes. The defaults of CYCLE=10, RETPD=60, LOWTHRESHOLD=50, and HIGHTHRESHOLD=80 will be used unless these have been changed in the FDR Global Option Table. Some of these volumes are SMS-managed, so a storage class is specified (it will be ignored for non-SMS volumes).

```
//STEP1 EXEC PGM=FDRABRM,REGION=256K
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
        ABRINIT VOLG=TEST,ENABLE=(SCRATCH,OLDBKUP),GEN=3,STORCLAS=SCTEST
        ABRINIT VOLG=PROD,ENABLE=OLDBKUP,GEN=6,STORCLAS=SCPROD
```

50.45 CONTINUED . . .

**INITIALIZE
SMS VOLUMES**

Set the ABR disk volume processing options for volumes starting with SMS, enabling all ABR features. Use the defaults for CYCLE, RETPD, GEN and LOWTHRESHOLD, but set HIGHTHRESHOLD to 90. Do not print the VTOCs The SMS ACS routines at this installation would not allow a data set with a name like that of the ABR Model DSCB to be placed on an SMS-managed volume, so it is necessary to specify BYPASSACS.

```
//STEP1      EXEC   PGM=FDRABRM,REGION=256K
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      ABRINIT      VOLG=SMS*,ENABLE=(OLDBKUP,SCRATCH),HIGHTHRESHOLD=90,
                   LOWTHRESHOLD=40,GEN=4,CYCLE=6,RETPD=30,
                   STORCLAS=ABRMODEL,BYPASSACS
```

**MODIFY ABR
OPTIONS**

Change the ABR options on disk volume PROD01, setting the tape retention period to 75 days and setting its low and high allocation thresholds to 60 and 90 percent. Change the ABR options on disk volume PROD02, prohibiting automatic archiving. Print changed DSCBs.

```
//MAINT      EXEC   PGM=FDRABRM,REGION=256K
//SYSPRINT    DD     SYSOUT=A
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      DEFAULT      LIST=YES
      MAINT         VOL=PROD01,RETPD=75,
                   LOWTHRESHOLD=60,HIGHTHRESHOLD=90
      MAINT         VOL=PROD02,DISABLE=ARCHIVE
/*
```

**REBUILD ABR
MODEL DSCB**

User has accidentally scratched the ABR Model DSCB. This job will reset the volume to the current GEN and CYCLE as stored in the ABR catalog. The default number of cycles for TYPE=AUTO is set to 6. The default retention period for the volume is set to 30 days, and the high allocation threshold is set to 95. The low threshold and the GEN= value will be set to the defaults in the FDR Global Option Table.

```
//REMODL     EXEC   PGM=FDRABRM
//SYSPRINT    DD     SYSOUT=A
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      REMODEL      VOL=TST001,CYCLE=6,RETPD=30,HIGHTHRESHOLD=95
/*
```

**CORRECT ABR
MODEL DSCB**

User has backleveled the ABR Model DSCB, perhaps by doing a restore with SAR (stand-alone Restore) or some other restore product. This job will correct the ABR Model DSCB to the current GEN and CYCLE.

```
//REMODL     EXEC   PGM=FDRABRM
//SYSPRINT    DD     SYSOUT=A
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      REMODEL      VOL=SYS001,FORCE
/*
```

50.46 FDRABRM ISPF INTERFACE

Most of the functions of FDRABRM, the ABR VTOC maintenance utility, can be executed using the ABR ISPF dialogs. The dialog can also be used to simply display the ABR status of online DASD volumes. Assuming that option A on your ISPF main menu gets you to the ABR dialogs, from the main menu enter: **A.I.8**

```

----- FDR INSTALLATION -- SET ABR DISK VOLUME PROCESSING OPTIONS -----
COMMAND ===>                                SCROLL ===> PAGE
PLEASE PRESS THE "ENTER" KEY TO DISPLAY THE TABLE

THE ABR DISK VOLUME PROCESSING OPTIONS MUST BE SET PRIOR TO ABR EXECUTION.

SETTING THE ABR PROCESSING OPTIONS DOES NOT AFFECT NORMAL USE OF THE VOLUME.

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME      ===> 'IDP.MODFDR53'
  VOLUME SERIAL       ===>

SYSOUT CLASS          ===> *

JOB STATEMENT INFORMATION:
  ===> //useridA JOB (ACCOUNT),'NAME',
  ===> //          NOTIFY=userid
  ===> // *
  ===> // *

```

ISPF PANEL**A.I.8**

The end result of this dialog will be to submit a batch job to execute FDRABRM with the proper control statements. On this panel, specify the name (and optional volume serial if not cataloged) of the FDR program library, the SYSOUT class to which output is to be directed (usually "**") and update the JOB statement displayed to be acceptable for your installation. When done, press ENTER to continue and display this panel:

```

----- FDR INSTALLATION -- SET ABR VOLUME PROCESSING OPTIONS Row 1 to 1 of 1
COMMAND ===>                                SCROLL ===> PAGE
PLEASE ENTER THE VOLUME SERIAL NUMBERS TO DISPLAY/SET ABR PROCESSING OPTIONS
  SUBMIT - SUBMIT FDRABRM BATCH JOB      EDIT - EDIT FDRABRM BATCH JOB
  FIND   - FIND THE SPECIFIED STRING     HELP - DISPLAY TUTORIAL PANELS

      VOLUME  ---CURRENT BACKUP---  MAX GEN  RETPD  ARCHIVE  STORCLAS  FORCE
      SERIAL      EXPDT      MAX AC  RETP2  SCRATCH  ARCHI
CMD  NUMBER  GEN CYC EXPD2 AC MOD      OLDBKUP  ARCLOW
-----
' ' ' test*                NO          4      60      YES                NO
                        10      0      NO          80
                        NO          NO          50
-----

```

The fields toward the right of the panel correspond to many of the operands of the ABRINIT, MAINT, and REMODEL statements; the values shown are the defaults in effect. The cursor is positioned in the "Volume Serial Number" column. To display the actual ABR status of one or more disk volumes, you can enter one of:

- a volume serial, such as PROD01
- a prefix followed by an asterisk (*), such as PROD*, to display all volumes starting with that prefix
- an asterisk (*) to display all online DASD volumes (this may be time consuming if you have a great many online volumes).

50.46 CONTINUED . . .

Press ENTER to display a panel such as:

----- FDR INSTALLATION -- SET ABR VOLUME PROCESSING OPTIO Row 1 to 4 of 4										
COMMAND ==>					SCROLL ==> PAGE					
SUBMIT - SUBMIT FDRABRM BATCH JOB					EDIT - EDIT FDRABRM BATCH JOB					
FIND - FIND THE SPECIFIED STRING					HELP - DISPLAY TUTORIAL PANELS					
	VOLUME	---CURRENT BACKUP---				MAX GEN	RETPD	ARCHIVE	STORCLAS	FORCE
	SERIAL	EXPDT				MAX AC	RETP2	SCRATCH	ARCHI	
CMD	NUMBER	GEN	CYC	EXPDT	AC	MOD		OLDBKUP	ARCLW	

	TEST01	****	NO	MODEL	**	NO	4	60	YES	NO
							10	0	NO	
								NO	80	
									50	

	TEST02	****	NO	BKUP	**	NO	4	60	YES	NO
							10	0	YES	
								YES	80	
									50	

	TEST03	447	1	96166	1	NO	5	35	NO	NO
							50	0	NO	
								YES		

						NO	4	60	YES	NO
							10	0	NO	
								NO	80	
									50	

This shows the current ABR status of the selected volumes. To add more volumes to the display, type another serial or prefix in the blank line at the bottom.

The CMD field will accept commands similar to ISPF EDIT: **I** or **In** will insert “n” blank lines where new serials can be entered. **R** or **Rn** will replicate the selected line “n” times. **D** or **Dn** will delete one or more lines.

In the fields headed “—Current Backup—”, the status of the most recent backups of the volume are shown; these are extracted from the ABR Model DSCB on the volume. The most recent generation and cycle numbers, the expiration dates of the most recent full-volume backups (EXPDT for COPY 1 and EXPD2 for COPY2 if it was created), and the current auto-cycle count are shown. If the volume has not been initialized for ABR processing (no ABR model DSCB in the VTOC), “NO MODEL” is shown. A volume which was initialized for ABR but which has never had a full-volume backup taken shows “NO BKUP”.

These fields are normally for reference only. But if they have somehow become inaccurate or back-leveled (such as a stand-alone restore (SAR) of the volume), you can overwrite them with the proper values. **However, if you do so you must also change the column labeled “MOD” to “YES”; otherwise the changes will not be accepted.**

The rest of the fields (except FORCE) represent the ABR initialization options for the volume. For an initialized volume, they are the actual values extracted from the ABR Model DSCB and you can overwrite any of them to change them. For an uninitialized volume they are the default values which you can accept or overwrite to change.

Enter YES in the FORCE column if an attempt to initialize or maintain a volume has failed (see the description of the FORCE operand under the various FDRABRM statements earlier in this section to see if FORCE is appropriate). Note that FORCE will be required if you have copied data sets from other ABR-managed volumes to this volume before initializing it for ABR.

When you type SUBMIT or EDIT on the command line, the dialog will generate an FDRABRM jobstream. It will include ABRINIT statements for any displayed uninitialized volumes, MAINT statements for any initialized volumes whose options you have changed, and REMODEL statements for any initialized volumes whose current backup information you have modified. EDIT gives you a chance to review and modify the jobstream before submission.

To exclude volumes from processing, simply delete them from the display (the **D** line command) before entering SUBMIT or EDIT.

50.50 CATALOG MAINTENANCE UTILITY (FDRABRCM)

FDRABRCM is an ABR utility program which performs maintenance on the ABR catalog and the ABR scratch catalog.

The ABR catalog is an ICF catalog which is used to store the location of ABR backup data sets. It is always used for ABR full and incremental backups, and is sometimes used for ARCHIVE backups.

The ABR scratch catalog is an ICF catalog used to store information on the ABR backups of DASD data sets which have been scratched (data sets which are still on DASD have their backup information in the DSCBs of the data sets themselves). Entries are built in the scratch catalog only if you have the ABR DADSM exit installed, which intercepts SCRATCH and RENAME requests.

In most installations, both catalogs are stored in the same ICF catalog, but they may also be separate. These are normal catalogs containing only non-VSAM catalog entries, although ABR uses some catalog fields for special purposes. You may manually delete entries with IDCAMS "DELETE NOSCRATCH" or IEHPROGM "UNCATLG" statements, but FDRABRCM will help automate this deletion.

FDRABRCM can:

- Uncatalog entries from ABR scratch catalog that have expired.
- Uncatalog entries from ABR scratch catalog that are within a specified date range.
- Uncatalog obsolete backup data set entries from the ABR catalog.
- Uncatalog backup data sets from the ABR catalog for volumes which no longer exist.
- Simulate all of the functions mentioned above.

**FDRABRCM
STATEMENTS**

FDRABRCM has only one statement, PURGE, with two variations:

PURGE SCRATCH - purges entries from the ABR scratch catalog.

PURGE BACKUP - purges backup entries from the ABR catalog

WHEN TO RUN

Under normal circumstances, the ABR catalog is self-maintaining. Entries for obsolete backup generations are automatically deleted by ABR as new generations are created; depending on your tape management system options and the retention type you use for ABR backups, your tape management system may also delete expired entries for tape backups. However, there are circumstances where obsolete backups might remain cataloged. A periodic execution of PURGE BACKUP with no other operands will insure that these entries are cleaned up.

Also when a DASD volume is no longer being backed up by ABR (such as a volume which no longer exists), a special PURGE BACKUP may be required to delete the backups when they are no longer required. Again, your tape management system may delete the entries when the backups expire.

The ABR scratch catalog is **not** self-maintaining. Entries in the scratch catalog point to backups in the ABR catalog; the scratch entries are obsolete when all of the backups they point to are deleted from the ABR catalog, but they do not go away by themselves. If you use the scratch catalog, you must periodically run PURGE SCRATCH with no other operands to identify and delete the obsolete entries.

Operands on each statement may be used to limit the operation to a subset of the catalog entries, or to specify other criteria for deletion of the entries.

50.51 FDRABRCM JCL REQUIREMENTS

The following Job Control Statements are required to execute FDRABRCM:

EXEC STATEMENT Must specify the name of the ABR catalog utility program (FDRABRCM). It may also contain the region requirements of 256K. The user can specify any 'PURGE' statement as PARM= data on the EXEC statement.

STEPLIB OR JOBLIB DD STATEMENT If required, must specify the load module library in which ABR resides. This library must be APF authorized.

SYSPRINT DD STATEMENT Specifies the output message data set. This is a required DD statement and is usually a SYSOUT data set.

SYSUDUMP DD STATEMENT Specifies the abend dump data set. Although not required, we strongly urge you to always include this DD statement so that we can help you diagnose your error conditions. Usually a SYSOUT data set.

If you have the ABEND-AID product from COMPUWARE include the following so that a fully-formatted dump is produced: `//ABNLIGNR DD DUMMY`

ABRMAP DD STATEMENT Specifies the optional report data set, usually a SYSOUT data set. When ABRMAP is not found within the JCL stream, the reports will be output to the SYSPRINT data set.

SYSIN DD STATEMENT Specifies the control statement data set, usually an input stream or DD * data set. It can be omitted if the control statement is specified by PARM= on the EXEC statement.

50.52 FDRABRCM PURGE SCRATCH STATEMENT

PURGE	SCRATCH	,LINECNT=nn
	SCR	,LIST=ALL ERR
	,BKDATE=yyddd yyyyddd	,SELTERR=YES NO
	,BKDAYS=nnnn	
	,DSN=dsname	,SIMULATE SIM
	,DSG=dsgroup	
	,FORMAT=PRT CRT	,VOL=vvvvvv
		,VOLG=vvvvv

**PURGE
SCRATCH
STATEMENT**

PURGE SCRATCH examines entries in the ABR scratch catalog and uncatalogs those that it considers to be obsolete. By default it will check all entries in the scratch catalog but you may limit it to checking only certain data sets by specifying the DSN= and/or DSG= operands. You may also limit it to checking only data sets that were scratched from certain DASD volumes with the VOL= or VOLG= operands.

Entries in the ABR scratch catalog point to one or more backup entries in the ABR catalog, defining the location of the backups of those scratched data sets. By default, PURGE SCRATCH will consider a scratch entry obsolete if all of the backups it points to are no longer cataloged, meaning that they have expired and are no longer available. If you specify the BKDATE= or BKDAYS= operand, PURGE SCRATCH will also purge scratch entries whose backups still exist as long as the most recent backup was taken before the date specified.

Unless you have special needs, no operands are usually required. Without other operands, PURGE SCRATCH will scan the entire scratch catalog and purge all entries whose backups have expired. PURGE SCRATCH should be run on a regular schedule but usually not more often than once a week.

OPERANDS

BKDATE= Specifies that any ABR backups pointed to by scratch catalog entries will be considered to be expired if they were created on or before the Julian date (year plus day number) specified; they will also be considered expired if they are no longer cataloged. The date may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are assumed to be in the 21st century, e.g., 12015=2012.015. For readability, a period may be inserted between the year and day (e.g., 1997.123).

BKDAYS= Specifies that any ABR backups pointed to by scratch catalog entries will be considered to be expired if they were created on or before the date calculated by subtracting BKDAYS= from today's date; they will also be considered expired if they are no longer cataloged.

If neither BKDATE= nor BKDAYS= is specified, ABR backups pointed to by scratch catalog entries are considered expired only if they are no longer cataloged.

DSN= Specifies one or more data set names (up to 44 characters in length). Multiple names may be specified in parentheses, separated by commas.

50.52 CONTINUED . . .

DSG=	<p>Specifies one or more data set name prefixes (up to 44 characters in length). Multiple prefixes may be specified in parentheses, separated by commas. . Only scratch catalog entries whose name begins with one of the prefixes will be examined.</p> <p>There is a special form of the DSG= operand. Leading periods (.) after DSG= indicates that the search argument begins after one (1) or more index levels. Each period indicates that one (1) or more index level is to be bypassed. For example, DSG=..ABC will select data sets whose third index level starts with ABC.</p> <p>The default is, if neither DSN= nor DSG= is specified, the data set name will not participate in entry selection. Both DSN= and DSG= may be specified on the same PURGE statement; each scratch catalog entry will be compared to the combined list. Since the catalog is a keyed file ordered by data set name, specifying DSN= or DSG= (except the special form) will speed up processing.</p>
FORMAT=	<p>Specifies the format of the reported generated by PURGE.</p> <p>PRT produces an 121 character line formatted for a printer, one data set per line.</p> <p>CRT produces an 80 character line suitable for display on a terminal, with multiple lines per data set.</p>
LINECNT=	<p>Specifies the maximum number of lines each report page can contain, from 10 to 99, inclusive.</p> <p>The default is that each page will contain a maximum of 58 lines.</p>
LIST=	<p>Specifies whether the scratch catalog entries selected for purge should be listed.</p> <p>ALL (the default) lists all selected entries.</p> <p>ERR lists only entries which were selected but not successfully purged (uncatalog error).</p>
SELTERR=	<p>If DSN=, DSG=, VOL=, or VOLG= were specified, but one or more of the values specified for those operands did not match any entries in the scratch catalog, SELTERR= controls whether this will cause FDRABRCM to end with a non-zero return code. Specify NO if you want the step to end normally even though there were some unmatched selection criteria.</p> <p>The default is YES.</p>
SIMULATE SIM	<p>Specifies that FDRABRCM is to report on all the scratch catalog names selected, but is not to actually purge any. This is used to verify that your PURGE statement will have the desired results.</p> <p>The default is PURGE processing will be attempted.</p>
VOL=	<p>Specifies a DASD volume serial number. Data sets found in the ABR scratch catalog which were scratched from that volume will be considered for purge. Multiple volume serial numbers may be specified in parentheses separated by commas.</p>

50.52 CONTINUED . . .

VOLG=

Specifies a DASD volume serial prefix. Data sets found in the ABR scratch catalog which were scratched from a volume beginning with this prefix will be considered for purge. Multiple prefixes may be specified in parentheses separated by commas.

The default is, if neither VOL= nor VOLG= is specified, entries will be considered regardless of the DASD volume from which they were scratched. Both VOL= and VOLG= may be specified on the same PURGE statement; each scratch catalog entry will be compared to the combined list.

50.53 FDRABRCM PURGE BACKUP STATEMENT

PURGE	BACKUP	,GEN=(nn,...,nn)
	,BKDATE=yyddd yyyyddd	,LINECNT=nn
	,BKDAYS=nnnn	,LIST=ALL ERR
	,COPY=n	,MAXGEN=nnnn DEFAULT
	,CYCLE=(nn,...,nn) ALL	,NOMODEL
	,CURRINFO=MODEL CATLG	,SELTERR=YES NO
	,FORCE	,SIMULATE
	,FORMAT=PRT CRT	,VOL=vvvvvv
		,VOLG=vvvvvv

**PURGE
BACKUP
STATEMENT**

PURGE BACKUP scans the ABR catalog for records of ABR backups that are obsolete and uncatalogs them.

For every DASD volume initialized for ABR processing, the ABR model DSCB in its VTOC contains the maximum number of generations (MAXGEN) that your installation wants to keep for that volume; the generation number is incremented for every full-volume backup. PURGE BACKUP takes the current generation number for each volume, minus its MAXGEN value; any backups for that volume still cataloged in the ABR catalog with generation numbers prior to that are eligible for purge.

So, with no other operands, PURGE BACKUP will uncatalog any ABR backups which are not within the range of generations that you have indicated should be kept for each DASD volume. Since normal ABR operation will also uncatalog these obsolete backups, this processing is required only to clean up backup entries that were not automatically uncataloged for some reason. Most installations should run PURGE BACKUP periodically (perhaps once a month).

Operands are available to override the MAXGEN value from the ABR Model DSCB, to limit the purge based on backup date, and to purge specific generations and/or cycles.

NOTE: FDRABRCM will never purge the current (most recent) generation of backup for any DASD volume unless you specify the FORCE operand. FORCE can be used to purge all generations of backups for a DASD volume that no longer exists.

OPERANDS**BKDATE=**

Specifies that ABR backups that are outside the MAXGEN limits will be uncataloged only if they were created on or before the Julian date (year plus day number) specified. The date may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are assumed to be in the 21st century, e.g., 12015=2012.015. For readability, a period may be inserted between the year and day (e.g., 1997.123).

50.53 CONTINUED . . .

- BKDAY=** Specifies that ABR backups that are outside the MAXGEN limits will be uncataloged only if they were created on or before the date calculated by subtracting BKDAY= from today's date.
- If neither BKDATE= nor BKDAY= is specified, ABR backups are uncataloged based only on the generation number (MAXGEN).
- COPY=** Specifies the copy number (1 or 2) you wish to uncatalog.
- The default is that both copies will be uncataloged.
- CYCLE=** Specifies one or more cycle numbers you wish to uncatalog. Only the specified cycles from within each generation selected will be purged. You may specify 'CYCLE=ALL', which will select all cycles in every selected generation **except** for cycle 00 (the full-volume backup) in each generation and the most recent cycle in the current generation (the most recent backup).
- The default is all cycles within selected generations, except the most recent cycle in the current generation (unless FORCE is also specified).
- CURRINFO=** Specifies what source contains the correct current generation and cycle values when the ABR Model DSCB and ABR backup entries conflict.
- MODEL** - the ABR Model DSCB contains the correct values
- CATLG** - the ABR backup entry with the highest backup date contains the most current information.
- The default is MODEL unless the DASD volume has no ABR model DSCB, in which case CATLG is used; in this case the MAXGEN= operand is also required since this information is stored only in the ABR Model DSCB.
- FORCE** Specifies that no checks or limitations are applied when selecting backups for uncatalog. This is usually used to completely uncatalog all backups for a volume which no longer exists or is no longer being backed up by ABR. This operand also **requires** the use of the **VOL=** operand.
- FORMAT=** Specifies the format of the reported generated by PURGE.
- PRT** produces a 121 character line formatted for a printer, one data set per line.
- CRT** produces an 80 character line suitable for display on a terminal, with multiple lines per data set.
- GEN=** Specifies one or more generation number(s) the user wishes to select for processing. This value can be any generation except the most current generation unless FORCE is also specified.
- The default, if GEN= is not specified on the PURGE statement, are the entries found outside the MAXGEN limits for the volume.
- LINECNT=** Specifies the maximum number of lines each report page can contain, from 10 to 99, inclusive.
- The default is that each page will contain a maximum of 58 lines.

50.53 CONTINUED . . .

LIST=	<p>Specifies whether the ABR backup catalog entries selected for purge should be listed.</p> <p>ALL (the default) lists all selected entries.</p> <p>ERR lists only entries which were selected but not successfully purged (uncatalog error).</p>
MAXGEN=	<p>Specifies a maximum number of generations to keep in the ABR backup catalog for each DASD volume. You may also code 'MAXGEN=DEFAULT' which directs the program to use the ABR default from the FDR Global Option Table, which is normally 4.</p> <p>The default is taken from the ABR model DSCB in the VTOC of each DASD volume.</p>
NOMODEL	<p>Specifies that only volume entries within the ABR backup catalog not having a corresponding ABR Model DSCB (missing from the VTOC or the volume is offline) will be selected for processing.</p> <p>The default is only ABR backup entries having an ABR Model DSCB available will be selected. When no model condition exists 'CURRINFO=CATLG' is set by the program which also requires 'MAXGEN=' operand to be specified.</p>
SELTERR=	<p>If VOL= or VOLG= were specified, but one or more of the values specified for those operands did not match any entries in the ABR catalog, SELTERR= controls whether this will cause FDRABRCM to end with a non-zero return code. Specify NO if you want the step to end normally even though there were some unmatched selection criteria.</p> <p>The default is YES.</p>
SIMULATE SIM	<p>Specifies that FDRABRCM is to report on all the ABR backup names selected, but is not to actually purge any. This is used to verify that your PURGE statement will have the desired results.</p> <p>The default is PURGE processing will be attempted.</p>
VOL=	<p>Specifies a DASD volume serial number. Only backups of that volume will be considered for purge. Multiple volume serial numbers may be specified in parentheses separated by commas.</p>
VOLG=	<p>Specifies a DASD volume serial prefix. Only backups of volumes beginning with this prefix will be considered for purge. Multiple prefixes may be specified in parentheses separated by commas.</p> <p>The default is, if neither VOL= nor VOLG= is specified, entries will be considered regardless of the DASD volume. Both VOL= and VOLG= may be specified on the same PURGE statement; each backup will be compared to the combined list</p>

50.54 FDRABRCM EXAMPLES

PURGE SCRATCH Scan the ABR scratch catalog and uncatalog all entries for scratched data sets whose backups no longer exist. This jobstream should be run periodically (typically once a week) by all ABR installations which use the scratch catalog (have the ABR DADSM exit installed) in order to cleanup the catalog.

```
//PURGE      EXEC  PGM=FDRABRCM
//SYSPRINT   DD    SYSOUT=A
//ABRMAP     DD    SYSOUT=A
//SYSUDUMP   DD    SYSOUT=A
//SYSIN      DD    *
PURGE        SCRATCH
```

PURGE SCRATCH BY DSGROUP Scan the ABR scratch catalog for scratched data sets beginning with one of several prefixes and uncatalog all entries whose backups no longer exist. This will execute in much less time than the previous example; it can be used to divide up the PURGE SCRATCH processing into smaller sets.

```
//PURGE      EXEC  PGM=FDRABRCM
//SYSPRINT   DD    SYSOUT=A
//ABRMAP     DD    SYSOUT=A
//SYSUDUMP   DD    SYSOUT=A
//SYSIN      DD    *
PURGE        SCRATCH,DSG=(TEST.,TSO)
```

PURGE BACKUP Scan the ABR catalog and uncatalog the ABR backups for any DASD volume if its generation number is outside the range of generations to be kept for the volume (the MAXGEN value in the ABR Model DSCB of the volume). This should be run periodically to cleanup obsolete backups which were not uncataloged by normal ABR processing for some reason.

```
//PURGE      EXEC  PGM=FDRABRCM
//SYSPRINT   DD    SYSOUT=A
//ABRMAP     DD    SYSOUT=A
//SYSUDUMP   DD    SYSOUT=A
//SYSIN      DD    *
PURGE        BACKUP
```

PURGE BACKUP WITH MAXGEN For several DASD volumes, you want to uncatalog some backups early, perhaps to free up tape volumes. Although the ABR Model DSCBs on each volume have a higher value for MAXGEN, this run will only keep the most recent 2 generations.

```
//PURGE      EXEC  PGM=FDRABRCM
//SYSPRINT   DD    SYSOUT=A
//ABRMAP     DD    SYSOUT=A
//SYSUDUMP   DD    SYSOUT=A
//SYSIN      DD    *
PURGE        BACKUP,VOL=(TEST01,TEST02),MAXGEN=2
```

PURGE BACKUPS OF OBSOLETE VOLUME When a DASD volume no longer exists or you no longer need to back it up with ABR, this example will result in the removal of all backup entries cataloged for the volume specified.

```
//PURGE      EXEC  PGM=FDRABRCM
//SYSPRINT   DD    SYSOUT=A
//ABRMAP     DD    SYSOUT=A
//SYSUDUMP   DD    SYSOUT=A
//SYSIN      DD    *
PURGE        BACKUP,VOL=LIB003,FORCE
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

50.70 FDRCLONE

FDRCLONE FDRCLONE is a separately-licensed enhancement to FDR/ABR.

FDRCLONE is a facility which "clones" disk volumes or data sets belonging to one MVS system to another MVS system, either an LPAR or a separate system. Its input is your normal FDRABR VOLUME backups, but data sets are actually restored only when they are needed. You can clone:

- All disk volumes in a data center
- selected disk volumes
- selected data sets

As cloned data sets are needed, they are dynamically restored from regular FDRABR volume backups (**full-volume ABR backups**) as described in the earlier parts of [Section 50](#). Only data sets which are actually needed by batch jobs or TSO users will be restored, so the total size of the restored data may be much less than the total in use at your home site.

FDRCLONE does not require any changes to batch JCL or TSO procedures. When a job or user references a cloned data set which has not yet been restored, the restore is automatically invoked.

This can be used to create test systems with copies of production data, such as a Y2K test system, yet may require far fewer disk volumes since only data sets which are actually used will be restored.

It can also be used in a disaster recovery environment. It is especially useful when the D/R site has fewer disk volumes or less total capacity than your production site, but it can also be used even when sufficient volumes are available. Many installations may require only 10-30% of their home site DASD capacity to restore the critical data sets necessary to run production.

FDRCLONE lets you identify the volumes and/or data sets to be "cloned". The catalog entries of the selected data sets will be specially marked for cloning. Those data sets will not be initially restored, but are dynamically restored when a batch job or TSO user references them. This allows testing (or, for a D/R site, production) to start much earlier, compared to waiting for all of the volume restores and/or application restores to complete.

On the clone system,

- you may choose to define disk volumes with the same volume serials in use at your home site, or a subset of them. In this case, no changes to SMS definitions or the control information of other allocation control software is required; however, each SMS storage group or allocation target group must have sufficient space to hold all the cloned data which is actually restored.
- Or you can use totally different volume serials, which may require updates to allocation controls.

FDRDRP FDRCLONE also includes FDRDRP, a utility for optimizing full-volume recovery from ABR volume backups (**full-volume and incremental backups**). This can be used for volumes where full-volume recovery is more appropriate, such as system volumes and high-priority production volumes which must be recovered quickly.

A normal ABR full-volume recovery restores one disk volume at a time. Since a given backup tape may contain the backups of many disk volumes, the same tapes may be unloaded and remounted over and over, taking considerable time and overwhelming tape librarians and automated tape libraries.

FDRDRP processes multiple full-volume recovery tasks in parallel. It manages usage of the backup tapes required for those restores, so that each backup tape is mounted a minimum number of times, **usually one mount per tape volume**. This will greatly reduce the elapsed time required to recover the volumes and eliminate most extra tape mounts.

FDRDRP is documented beginning in [Section 50.80](#).

50.70 CONTINUED . . .

**FDRCLONE
TARGET
SYSTEMS**

The system on which you execute FDRCLONE will usually be one of these:

- an LPAR on the same system which executes your production work. This LPAR will usually have access to all of the production disk volumes via shared DASD (although they may be marked offline to MVS on the LPAR) as well as the test volumes designated for its use.
- a separate processor in the same complex as the system which executes your production work. This processor might have access to all of your production volumes via shared DASD, or it might be totally isolated with its own unique disk volumes.
- A totally separate system, isolated from the production system with its own disk volumes. This might be a Disaster/Recovery site. It is probably physically distant from the production system.

On the clone target system, volumes to be restored are generally divided into these categories:

- system volumes (volumes necessary to IPL and run the system), such as system residence, spool, paging, common libraries and catalogs containing those data sets. These volumes should be restored from full-volume backups with FDR, ABR or SAR so that the target system can be IPLed
- production data base volumes such as CICS data volumes. It is often more efficient to restore them with full-volume restores or explicit data set restores
- other production application data volumes
- test data volumes
- TSO volumes

You may choose to restore data in the last 3 categories can be restored with manual FDR or ABR full-volume. The FDRDRP program, a part of FDRCLONE, can speed up full-volume recovery from ABR backups when such restores are necessary. It is documented in [Section 50.80](#).

50.70 CONTINUED . . .

**THE
FDRCLONE
SOLUTION**

FDRCLONE addresses all of these problems and limitations by allowing application and TSO data sets to be restored from your normal ABR volume backups only when they are needed. You only need to restore the system volumes and possibly production data bases or other data sets which are known to be in frequent use (such as some common libraries). Testing or resumption of production work can begin immediately.

The FDRCLONE process is similar to that used by FDRABR Auto-Recall from Archive backups. However, data is restored from normal FDRABR volume backups.

To use FDRCLONE you must have these resources at the test or D/R site:

- a running system (system volumes).
- a set of ABR full-volume backups for all the original volumes to be cloned. This might be your normal, scheduled backups, or might be special backups taken to provide restores to a particular point-in-time. If you use CLONE TYPE=DEFER, you will also need the incremental backups for those volumes unless you deliberately restrict the recovery to full-volume backups.
- a backup of the ABR catalog and all user catalogs which contain data sets which will be cloned. This backup must be taken immediately after the ABR backups (above) have completed.
- If you are executing CLONE TYPE=DEFER at the original site, that execution produces a sequential file which must be taken to the test/DR site. This contains a record of all catalog updates necessary to mark data sets as "cloned" on the clone target system.

At the target test or D/R system, you will run program FDRCLONE which selects the volumes and or data sets you want to clone and marks their catalog entries with special indicators. At this point, the "cloned" data sets have not been restored but their catalog entries indicate that they can be restored.

You must activate the ABR Catalog Locate exit (from FDRABR V5.3 level 30 or above) if you don't already have it active. **The FDRCLONE option in the FDR option table must be enabled before you activate the exit.**

The first time that a data set marked for FDRCLONE processing (a "cloned" data set) is referenced by a batch job or TSO user, it will be dynamically restored. The ABR Catalog Locate exit will recognize each cloned data set and invoke ABR to do the restore.

So, the first reference to a cloned data set will automatically mount the required ABR backup tapes and restore the data set. The target volume for each restore is chosen by the normal ABR rules. The target may be chosen by SMS or by other allocation control software you have installed. If you have implemented the ABR Restore Allocation List, it may be used to select the target, or ABR may attempt to restore each data set to its original volume serial (if available). The target volumes must have sufficient space to hold all of the restored data sets.

If several cloned data sets are referenced by multiple jobs, but reside on the same backup tapes, FDRCLONE will attempt to restore them in one pass of the backup tape. Even if a restore request requiring a tape comes in after that tape is mounted, FDRCLONE will pass the tape to the next restore task which needs it without dismounting it.

50.70 CONTINUED . . .

**BENEFITS
ON TEST
SYSTEMS**

On a Y2K test system, or other test systems which use copies of production data, FDRCLONE will:

- simplify the creation of the test data. There is no need to do special backups or analyze what data sets are used by each application system to be tested. Data sets required by tests will be automatically restored.
- allow testing to begin almost immediately.
- reduce the amount of disk space devoted to the test system, since only data sets actually required by tests will be restored.
- reset the test system. If you need to repeatedly run tests starting with data from the same point-in-time, you only need to reinitialize the clone target volumes to delete all of the cloned data sets from disk (if you like, you can use ICKDSF INIT to quickly delete the data sets), re-restore the ABR and user catalogs and re-execute the FDRCLONE job. FDRCLONE will clone the data sets from the same backups as before, giving you the same testing environment used previously.

**BENEFITS
FOR
DISASTER/
RECOVERY**

At a disaster/recovery site, during a test or a real disaster, FDRCLONE will:

- simplify the recovery process. It is no longer necessary to do full-volume recoveries or application backup restores before resuming processing. Data sets required by jobs or TSO users will be automatically restored.
- allow testing or resumption of production to begin soon after an IPLable system has been created.
- reduce the amount of disk space required on the D/R system, since only data sets required by production jobs and TSO users will be restored. Since D/R sites typically charge based on the amount of disk space you need, this may result in a real cost savings.

When a real disaster is over, and you are returning to your re-activated production system, you need to carry back full-volume backups of all the volumes which were targets for the cloned data sets (usually far fewer volumes than those at your production site). If your production disk volumes survived the disaster, you can just do data set restores from those backups; FDR will locate the older version of each data set and restore the updated data on top of it. If you are restoring to completely new volumes, you need to recover each volume from the same ABR backups that were current at the time of the disaster (those read by FDRCLONE), and then do the data set restores from the backups taken just before you left the disaster site.

50.71 FDRCLONE OPERATION

There are two basic ways of executing FDRCLONE:

- **CLONE TYPE=INPLACE** - this is executed only on the clone target system. It will clone data sets only from ABR full-volume backups, since it has no way to determine which incremental backups contain backups of individual data sets. If restoring data sets from the latest full-volume backup is acceptable, this is the simplest procedure. If you are cloning a system for testing, rather than for disaster/recovery, this presumes that you took ABR full-volume backups of the original system at the point-in-time that you wanted to "clone".
- **CLONE TYPE=DEFER and TYPE=ACTIVATE** - CLONE TYPE=DEFER is executed on the original system. Since the original data sets are still available, FDRCLONE can determine the most recent ABR generation and cycle of the backup of each data set so it can clone each data set from the appropriate backup. It makes no changes to the running system, but creates a data file which is used on the clone target system.

CLONE TYPE=ACTIVATE is executed on the clone target system. It reads the data file produced by TYPE=DEFER and makes the appropriate updates to the cloned catalogs.

IMPORTANT WARNING: on V5.3 level 50, FDRCLONE only supports CLONE TYPE=INPLACE. Support for TYPE=DEFER/ACTIVATE will be provided in a future release.

**ON YOUR
PRODUCTION
SYSTEM**

- **ABR BACKUPS:** create ABR volume backups for each disk volume containing data sets to be cloned. If you are going to a disaster/recovery site, these will be your normal periodic ABR full-volume backups. For a test system, you may want to execute special ABR backups to create a well-defined point-in-time image of your data; you can even repeatedly reset the test system to that point. These backups must be available at the clone target system; for disaster/recovery you will want to send them to offsite storage. FDR InstantBackup, a separately licensed component of FDR documented in [Sections 25-29](#), may be able to help create the point-in-time backups with minimal disruption.
- **FDRCLONE: immediately before the catalog backups**, if you intend to use the CLONE TYPE=DEFER and TYPE=ACTIVATE procedure, execute FDRCLONE with TYPE=DEFER, directing TAPE1 to a data set on tape. This data set contains "before and after" images of the catalog records which will be updated on the clone target system. This file must be available at the clone target system; for disaster/recovery you will want to send it to offsite storage. See example jobstreams in [Section 50.75](#).
- **CATALOG BACKUPS: Immediately after the FDRCLONE execution and ABR full-volume backups are complete**, do a separate FDRAPPL (application) backup or FDRDSF backup of the ABR catalog plus any other user catalogs which contain data sets which will be cloned (actually you will probably want to include all user catalogs). This should be done immediately after the ABR backups are complete, so that the catalogs will accurately reflect the contents of the volumes just backed up. These backups must be available at the clone target system; for disaster/recovery you will want to send them to offsite storage.

50.71 CONTINUED . . .

**ON THE
TEST/DR
SYSTEM**

- **IPL SYSTEM:** If necessary, use FDR, ABR, or SAR to restore the system volumes necessary to IPL the system. Then IPL it. See “Disaster/Recovery” in [Section 80](#) for more details.
- **CATALOG RESTORES:** restore the ABR catalog and the user catalogs backed up at your production site. Note that when the backup and restore are both done with V5.3 level 30 and above, a restore of a non-existent catalog will define the catalog, connect it to the master catalog and also define all of its aliases in the master catalog, so it is immediately usable.
- **INSTALL FDR:** If you do not have FDR installed on the test system, you must install it. You can install it from the current FDR distribution tape, but it is easier to restore a copy of the current FDR program library from your production system. If you do a complete re-install, you will have to make sure that all FDR options in the FDR program library are the same as those used on the production system. [Section 91](#) has complete details on installation and options. The activation of FDRCLONE will be simplest if the FDR program library is in the system link list.
- **MARK DATA SETS FOR CLONING:** execute FDRCLONE to mark the catalog entries of all of the selected data sets for cloning.
 - 1) if you executed CLONE TYPE=DEFER on the original system, execute CLONE TYPE=ACTIVATE here. This will read the catalog update file created by the TYPE=DEFER step and will clone each data set from its most recent backup as of the time the TYPE=DEFER step was run.
 - 2) if you did not execute CLONE TYPE=DEFER, execute CLONE TYPE=INPLACE here. This will clone each selected data set from the most recent ABR full-volume backup (CYCLE=00) of the volume it is cataloged to, as recorded in the ABR catalog that you just restored.
- **ENABLE FDRCLONE:** Enable FDRCLONE restores as described below.
- **DEFINE RESTORE TARGET VOLUMES:** you have two options for defining the target volumes for the restore of cloned data sets:
 - 1) **if your clone target system contains disk volumes identical to your original system** (same number, size and volume serials), this step is unnecessary. FDRCLONE will use the original volume serial of every data set as its target unless overridden by SMS or other allocation control software.
 - 2) **if the clone target system has fewer volumes or different volume serials**, you must define where FDRCLONE is to restore data sets. This will usually be required for FDRCLONE.
 - a) if the cloned data sets are SMS-managed, you may have to update your SMS storage group definitions and possibly your ACS routines to direct the restored cloned data sets to the proper volumes.
 - b) if you have third-party allocation control software (such as POOL-DASD, ACC, PRO-SMS or SAMS:ALLOCATE), you may have to update the allocation rules to direct the restored cloned data sets to the proper volume.
 - c) if the cloned data sets are not managed by SMS or a third-party allocation product, you must update the ABR Restore Allocation List on the clone target system to specify new target volumes for data sets from your original system. **The Restore Allocation List is described in detail in [Section 90.32](#)**, including instructions for updating it using the FDR ISPF Install dialogs. Here is a sample:

50.71 CONTINUED . . .

```

FDR INSTALLATION -- RESTORE ALLOCATION LIST ----- Row 1 to 4 COMMAND ==>
SCROLL ==> PAGE

SAVE   - SAVE CHANGES      COPY - COPY THE MODULE      CANCEL - EXIT IMMEDIATELY
RESET  - RESET MODULE       FIND  - FIND A STRING      HELP   - TUTORIAL PANELS

CMD   ALLOCATE STATEMENT (NOTE: VALIDATION IS DELAYED UNTIL MODULE IS SAVED)
---   -----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-
      ALLOCATE ALLDSN,VOL=TEST23,NVOL=TEST25
      ALLOCATE ALLDSN,VOLG=PROD,NVOLG=PROD
      ALLOCATE DSN=SYS2.***,VOL=(SYS123,SYS124),NVOL=SYS456
      ALLOCATE DSN=SYS3.***,NVOLG=SYS*
      ALLOCATE ALLDSN,NVOL=(POOL1*,POOL2*)
      ALLOCATE ALLDSN,NVOL=(POOL1*,POOL2*)

```

Any data set from TEST23 will be restored to TEST25.

Any data set from a volume starting with PROD will be restored to any current volume starting with PROD.

Any SYS2 data set from volumes SYS123 or SYS124 will be restored to SYS456.

Any SYS3 data set will be restore to any current volume starting with SYS.

All other data sets will be restored to current volumes beginning with POOL1 or POOL2.

- **RUN JOBS:** Start your testing or resume your production.

ENABLING FDRCLONE

On the clone target system, you must enable FDRCLONE restores before running any jobs which need the cloned data sets. Panel names in this section refer to installation panels in the FDR ISPF interface; they are described in more detail in [Section 90](#).

- 1) If you use ABR Archiving and auto-recall on your production system, you already have the ABR Catalog Locate Exit enabled so skip to step (5). If not, you must enable it.
- 2) Enable dynamic installation of the Locate Exit on panel A.I.4.11.1 as shown here:

```

----- FDR INSTALLATION -- SET DYNAMIC INSTALLATION EXITS -----
COMMAND ==>
ABRLOC      DYNAMICALLY INSTALL THE ABR CATALOG LOCATE EXIT..... YES

```

- 3) Some of the options for auto-recall of archived data sets also apply to FDRCLONE. They are on panels A.I.4.11.2 and A.I.4.11.3 as shown below:

```

FDR INSTALLATION -- SET ABR GLOBAL AUTO-RECALL OPTIONS -----
COMMAND ==>
LXFOREST  USER IS ASKED TO CONFIRM RESTORE..... YES
LXDREST   RESTORE TYPE(S) FOR TSO AUTO-RECALL..... FG,BG,RQ
          ( NO NONE FG FOREGROUND BG BACKGROUND RQ REMOTE QUEUE )
LXFGSYNBG CONVERT TSO FG RESTORE TO SYNCHRONOUS BACKGROUND RESTORE.. NO
LXFGERR   RETRY AUTO-RECALL REQUEST IF FG FAILS ( NO, BG OR RQ )... NO
LXNOMSG   SUPPRESS AUTO-RECALL MESSAGES IF TSO RECALL IS DISABLED... NO
LXNEWVOL  USER MAY SPECIFY A NEW VOLUME SERIAL WHERE TO RESTORE.... YES
LXREISSU  PROMPT TO CONFIRM THE SPECIFIED NEW VOLUME SERIAL..... NO
LXSYNPROC CATALOGED PROCEDURE NAME USED FOR BACKGROUND RESTORES.... SYNRECAL
LXCHKSEC  AUTO-RECALL SECURITY OPTIONS
LXBYPSEC  CHECK SECURITY BEFORE THE AUTO-RECALL RESTORE OPERATION... NO
          BYPASS SECURITY DURING AUTO-RECALL RESTORE OPERATION..... NO

```

50.71 CONTINUED . . .

:

```

----- FDR INSTALLATION -- MORE ABR GLOBAL AUTO-RECALL OPTIONS -----
COMMAND ==>
LXNCDENY   FAIL LOCATE WITH 'NOT CATALOGED' IF TSO USER DENIES RECALL YES
LXCONUSE   USE CONSTANT NEW VOLUME FOR ALL DATA SET RECALLS..... NO
LXCONVOL   CONSTANT NEW VOLUME FOR ALL DATA SET RECALLS..... NEWVOL
LXALTMSG    ISSUE ALTERNATE FDRW71 (CONFIRM RESTORE) MESSAGE FORMAT... NO
LXUNCAT    ASK USER WHETHER TO UNCATALOG INSTEAD OF RECALLING DATASET NO
LXMAXSTC   MAXIMUM NUMBER OF RECALL STARTED TASKS ACTIVE AT ONE TIME.. NONE
LXMAXREC   LIMIT OF TOTAL RECALLS (INCLUDING STC) ACTIVE AT ONE TIME.. NONE
LXDIRVOL   DIRECTED VOLUME SERIAL NUMBER TO BE RETURNED TO JES3..... NONE
LXDIRTYP   DIRECTED DEVICE TYPE TO BE RETURNED TO JES3..... NONE

```

The most important option is LXSYNPROC. In step (4) you will be installing a cataloged procedure used for starting the FDRCLONE restore task. If the name SYNRECAL is not acceptable, change it here. You can review the meanings of other options in [Section 90.22](#), but in most cases they will not need to be changed.

- 4) The FDR Installation Control Library (ICL) contains a member SYNRECAL. This member must be copied to a cataloged procedure library accessible by JES on your clone target system. If you changed the name on the A.I.4.11.2 panel, rename it to that name. If the library containing FDRCLONE is not in the system linklist, you will have to add a STEPLIB to the SYNRECAL proc.

- 5) Enable FDRCLONE restores on panel A.I.4.13

```

FDR INSTALLATION -- SET FDRCLONE OPTIONS -----
COMMAND ==>
FDRCLONE   FDRCLONE RESTORE IS ENABLED..... YES
MAXTAPES   MAXIMUM NUMBER OF TAPE UNITS FOR FDRCLONE RESTORES..... 10

```

you may also want to change the maximum number of tape drives to be used during FDRCLONE restores.

- 6) Dynamically install the exit and options with JCL similar to:

```

//FDRSTART EXEC PGM=FDRSTART
//SYSUDUMP DD SYSOUT=*

```

If your FDR program library is not in the system linklist, you will have to add these DDs:

```

//STEPLIB DD DISP=SHR,DSN=fdr library
//SYSLIB DD DISP=SHR,DSN=fdr library

```

FDRSTART will install the Catalog Locate Exit if you don't already have it installed, and will install or refresh the FDR options to enable FDRCLONE without requiring an IPL. However, note that the Locate Exit must be from ABR V5.3 level 30 or above to include the FDRCLONE support. When a job or user includes JCL referencing a cataloged data set, or does a dynamic allocation of a cataloged data set, the Locate Exit will identify data sets marked as clones and will automatically invoke FDRABR to restore them. Note that if you use ABR Archive with auto-recall on your production system, any data sets which were archived at the time that you cloned the system will also be recalled normally from archive if they are referenced on the clone system.

50.71 CONTINUED . . .

**AUTO-
MATIC
DATA SET
RESTORE**

Whenever a batch job, started task, or TSO user references a cataloged data set which has been marked as “cloned”, the ABR Catalog Locate Exit will detect this and initiate the restore of the data set. The process is similar to that used for auto-recall of archived data sets (described in more detail in [Section 51](#)):

- when a batch job or started task references one or more cloned data sets in a step, **all** cloned data sets referenced in that step will be restored at once. The job or task will be suspended until the restores are complete, so they will not be aware that cloning took place. This scenario is also true for data sets referenced by the logon proc of a TSO user.
- When a batch job or started task dynamically allocates a cloned data set, that data set will be restored by itself. The job or task will be suspended until the restore is complete. In the case of a multi-tasking program, such as DB2, only the task requesting the allocation is suspended but allocation ENQs may prevent other tasks from allocating or opening data sets.
- When a TSO user dynamically allocates a data set (almost all data set references from TSO result in a dynamic allocation), the user will be informed that a cloned data set has been referenced and given the option of restoring it immediately (foreground), or restoring it asynchronously (background) which allows the user to do other work while the restore takes place. You may want to customize the ABR TSO auto-recall option so that such clone restores are done automatically without waiting for a user response; [see Section 51.34](#)
- There are options in the FDR Global Option Table which restrict the maximum number of auto-recalls which can take place at once. Those options also limit the number of concurrent FDRCLONE restore tasks (but not the number of data sets which can be recalled); if a new request causes limits to be exceeded, the requesting jobs are suspended until other restores complete.

But there are some differences from auto-recall:

- All FDRCLONE restores are done in external started tasks, unlike auto-recall where an external task is only used for certain recalls. The name of the external task and the JCL proc which invokes it is specified by the SYNRECAL option in the FDR Global Option Table. The SYNRECAL proc you use for auto-recall will also work for FDRCLONE.
- When a restore from a given backup is complete, FDRCLONE will check to see if any new requests for the same backup tape have been queued during that restore. If so, they will be processed immediately.
- FDRCLONE manages the restores, and may start more than one restore task if a job references data sets whose backups reside on different backup tape volumes, in order to process the restores in parallel. The FDR Global Option Table contains an option to limit the number of tape drives which can be in use concurrently by FDRCLONE (MAXTAPES=).
- For batch jobs and started tasks, messages relating to the restore (such as FDR311 messages documenting the successful restore) will be displayed in your system log (SYSLOG) with the ID of the job or task but will not be printed in the joblog of that job or task. Except for the possible time delay due to the restore, jobs and started tasks will not be aware that the data sets were cloned and restored. TSO users will receive the messages when they reference cloned data sets.
- If a data set is restored from a particular backup tape by FDRCLONE, there is a high probability that the tape will be required for another data set shortly after. To avoid mounting the tape twice, FDRCLONE will leave the tape mounted when it is done with it, so that another FDRCLONE restore can use it without remounting it. The tape volume will be marked “de-mountable” so that if another job (even another FDRCLONE restore) requires a different tape volume and no other drives are free, MVS will dismount the tape and reuse the drive. If necessary, issue the MVS UNLOAD(U) console command to dismount the tape.

50.71 CONTINUED . . .

**MULTI-
VOLUME
DATA SETS**

If multi-volume data sets are to be cloned, and you are cloning by volume, e.g.,

```
SELECT CATDSN=**,VOL=(xxx*,yyy*)
```

which is recommended, instead of by data set

```
SELECT CATDSN=ABC**
```

then you must insure that all volumes which may contain pieces of a multi-volume data set are processed by FDRCLONE. Otherwise the data set will not be completely restored and will not be usable.

It is possible that you could use full-volume recovery for volumes containing part of a multi-volume data set and FDRCLONE for the rest. But if you do you must insure that all the pieces of the data set are from the same point-in-time. Otherwise the restored data set may contain inconsistent or incorrect data.

**RE-CLONING A
SYSTEM**

On a test system, you may want to re-clone the original data, setting all application data sets back to their original contents so that you can rerun the test jobs. On a Y2K system, this might be done so that you can test with the system date set to various values.

To re-clone:

- Delete all of the cloned data sets from disk. The easiest way to do this is to simply reinitialize the clone target volumes with ICKDSF. If there are other, non-cloned data sets on those disks, you will have to delete the cloned data sets individually; FDRABR SUPERSCRATCH (see [Section 51](#)) can be used to do this.
- Re-restore all of the catalogs (ABR catalog and user catalogs).
- Re-run FDRCLONE.
- Start your testing again.

On a test system, you may want to re-clone but use backups from a different point in time. The procedure is the same, but you must restore the catalogs corresponding to that new point-in-time; this will cause FDRCLONE to use the new ABR backup tapes as input.

50.72 FDRCLONE JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform FDRCLONE functions:

**STEPLIB or
JOB LIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRCLONE resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRCLONE), region requirement (REGION=4M or more), and optional PARM= operand.

If a PARM field is specified, FDRCLONE will use data specified as the first control statement, which must be a valid CLONE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//CLONE EXEC PGM=FDRCLONE,PARM='CLONE TYPE=INPLACE/ SELECT CATDSN=**,VOL=ABC*'
```

If FDRCLONE is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**SYSPRINT
DD
STATEMENT**

Specifies the primary output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**TAPE1 DD
STATEMENT**

Specifies a data set which contains "before and after" copies of catalog records modified by FDRCLONE. It is a normal sequential QSAM file DCB characteristics RECFM=VB, LRECL=32756, BLKSIZE=32760. It may be on tape or disk, but if FDRCLONE is being used for disaster/recovery, you will want to write it to tape and send it offsite with other backup tapes.

For CLONE TYPE=DEFER

TAPE1 is an output file, so you will want to specify the JCL parameters DSN=, UNIT=, and DISP=(,CATLG). DCB parameters will be ignored. You may need to specify tape retention information. FDRCLONE will write to TAPE1 the original catalog record and the modified catalog record for each selected data set, but the catalog will not be modified. This file will be read when you execute CLONE TYPE=ACTIVATE or TYPE=DECLONE so it must be available at the clone site along with the ABR backups current at the time that FDRCLONE was run. For example,

```
//TAPE1 DD DSN=FDRSYS.CLONE.SYS1(+1),UNIT=CART,DISP=(,CATLG),LABEL=RETPD=60
```

For CLONE TYPE=INPLACE

TAPE1 is an optional output file. However, it is required at a disaster/recovery site if you intend to eventually move your production work back to your original or rebuilt data center.

If used, it will look like the TAPE1 DD described above for CLONE TYPE=DEFER and will have the same contents. However, the catalog is updated immediately by TYPE=INPLACE, so this serves as a record of the catalog changes that were made and can also be used as input for CLONE TYPE=DECLONE.

If you don't want to create this file, TAPE1 must be a DUMMY data set, e.g.,

```
//TAPE1 DD DUMMY
```

For CLONE TYPE=ACTIVATE

TAPE1 is an input file created by a CLONE TYPE=DEFER step. FDRCLONE will read it and update your system catalogs with the changes that were recorded in the file.

For CLONE TYPE=DECLONE

TAPE1 is an input file created by a CLONE TYPE=DEFER or TYPE=INPLACE step. FDRCLONE will read it and update your system catalogs by reversing the changes that were recorded in the file.

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for ABR. Usually a DD * data set.

It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

50.73 FDRCLONE CLONE STATEMENT

CLONE TYPE=INPLACE|DEFER|ACTIVATE|DECLONE
 ,COPY=n
 ,MAXCARDS=nnnn

**CLONE
STATEMENT**

The CLONE statement initiates an update of your system catalogs to mark the selected data sets as being "cloned". Only the catalog entries of the selected data sets are modified by this program.

For TYPE=INPLACE and TYPE=ACTIVATE, the catalogs containing the selected data sets must be restored and connected to the master catalog of the running system with the proper aliases before FDRCLONE is executed. Note that FDR automatically connects user catalogs to the master catalog if the output catalog must be allocated (does not already exist). Also, if the backup and restore of catalogs are done with FDR V5.3 level 30 or above, the restore will automatically reestablish the aliases of each allocated catalog.

OPERANDS TYPE=

Specifies the type of FDRCLONE operation. [See Section 50.71](#) for more details. It is required.

INPLACE - performs a "clone-in-place" operation, immediately updating the system catalogs of the system on which it is executing. This should be executed only on a clone target system. Note that it will only clone data sets from ABR full-volume backups (cycle 00) and will not use incremental backups (cycle>00).

DEFER - creates a data file on TAPE1 which identifies all of the catalog changes necessary to clone the selected volumes and/or data sets on a clone target system. It does not update the catalogs. It should be run on the original system which is being cloned. It will clone individual data sets to their most recent backup, including incrementals (CYCLE>00).

ACTIVATE - reads the TAPE1 data file created by a TYPE=DEFER step and updates the system catalogs using the records in that data file. This should be executed only on a clone target system. This step should have the same SELECT statements (or a subset of them) as the TYPE=DEFER step which created the data file.

DECLONE - reads the TAPE1 data file created by a TYPE=DEFER or TYPE=INPLACE step and updates the system catalogs to backout some of the catalog changes made by that step. This is executed on the original "cloned" system. It is used when you return from a disaster/recovery site.

IMPORTANT WARNING: in FDRCLONE V5.3 level 50, only TYPE=INPLACE is supported. Other TYPE= options will not work. Support for TYPE=DEFER, ACTIVATE and DECLONE will be provided in a future release.

50.73 CONTINUED . . .

COPY=

Specifies the copy of the backup from which the restore of the cloned data sets is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRSEL utility ([See Section 60](#)). COPY= is ignored for CLONE TYPE=ACTIVATE and DECLONE.

If COPY=1 or 2 and FDRCLONE finds that one of the backup tapes is not cataloged under the copy specified, it will check to see if the other copy was created. If cataloged, FDRCLONE will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) FDRCLONE will automatically use the other copy (1 or 2) if it still exists.

Default is COPY=1 unless overridden by option BKPCOPY in the FDR Global Option table ([See Section 90](#)). For CLONE TYPE=DEFER, the option table in effect at the time the TYPE=DEFER job is run will provide the default.

MAXCARDS=

Enables FDRCLONE to accept additional SELECT and EXCLUDE statements during this execution. You can specify values up to 65535, but values over 100 will increase the region required by FDRCLONE.

Default is 100 statements.

Note: a large number of SELECT statements is not recommended for FDRCLONE. Each SELECT statement will cause one or more of your system catalogs to be scanned for matching data sets, which may dramatically increase the elapsed time of the FDRCLONE job. Rather than listing volumes to be cloned one at a time, e.g.,

```
SELECT CATDSN=**,VOL=PROD01
SELECT CATDSN=**,VOL=PROD02
```

It is better to list them on one statement:

```
SELECT CATDSN=**,VOL=(PROD01,PROD02)
or
SELECT CATDSN=**,VOL=PROD*
```


50.74 FDRCLONE SELECT AND EXCLUDE STATEMENTS

SELECT	CATDSN=filter
S	
	,STORGRP=(sgroup1,sgroup2,...)
EXCLUDE	,VOL=(vvvvv1,vvvvv2,...)
X	
	,CATALOG=catname
	,CATLIMITGDG=n
	,COPY=n
	,GEN=gggg
	,PRTALIAS

**SELECT/
EXCLUDE
STATEMENTS**

These statements select cataloged data sets to be marked as "cloned". FDRCLONE will search the system catalogs for catalog entries which match the statements you have provided and will mark the catalog entries so that the ABR Catalog Locate Exit will cause them to be automatically restored if any batch job or TSO user references them.

A **SELECT** statement identifies data sets or volumes to be processed. FDRCLONE will find data sets matching this specification in your system catalogs and mark them as cloned.

An **EXCLUDE** statement identifies data sets from within those selected by SELECT statements which are not to be processed.

All data sets in your system catalogs will be compared to these statements to identify those to be processed; each data set will be compared to each control statement until a match is found. A maximum of 100 of these control statements may be used in one execution unless overridden by MAXCARDS=.

If the first index level of CATDSN= is unqualified (no special selection characters,e.g., CATDSN=ABR.X**) only the system catalogs with matching aliases will be searched.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since FDRCLONE will only clone data sets which are selected, EXCLUDE statements can be used to exclude certain data sets from within a larger group on a SELECT statement

Note: a CLONE TYPE=DEFER step which creates a FDRCLONE data file, and the CLONE TYPE=ACTIVATE step which processes that data file should contain the same SELECT/EXCLUDE statements. The ACTIVATE step will search for the same catalog entries and will generate an error if an updated catalog record is not found in the data file for a selected catalog entry. However, the SELECTs in the ACTIVATE step may select a subset of the data sets which were selected by the DEFER step; only those data sets selected will be marked "cloned".

50.74 CONTINUED . . .

**PERFORMANCE
NOTE**

For CLONE TYPE=INPLACE and TYPE=DEFER, each SELECT statement will initiate a catalog search. To get the best performance, you need to select as many data sets on each SELECT as possible.

If you had these 3 SELECT statements:

```
SELECT CATDSN=**,VOL=ABC123
SELECT CATDSN=**,VOL=ABC234
SELECT CATDSN=**,VOL=ABC345
```

each of these will search ALL of your system catalogs separately, looking for data sets cataloged to the given volume. However,

```
SELECT CATDSN=**,VOL=(ABC123,ABC234,ABC345)
```

Will search the catalogs only once for all 3 volumes. You can place up to 255 volumes on a SELECT statement (continue the statement after any comma), e.g.,

```
SELECT CATDSN=**,VOL=(ABC123,ABC234,ABC345,TSO*,
PAY*,PROD*)
```

For disaster/recovery, Innovation recommends cloning by volume using CATDSN=** and VOL= or STORGRP=.

For testing (including Y2K testing), you may want to limit cloning to certain data sets with CATDSN=mask or CATALOG=.

OPERANDS**CATDSN=**

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80](#). It is required.

If the VOL= or STORGRP= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected. If multi-volume data sets are being cloned, you must insure that all volumes on which the data sets reside are included in the clone process.

```
Examples: CATDSN=USER1.JCL.CNTL
          CATDSN=**MASTER(0)
          CATDSN=PROD++.**.LIB*
          CATDSN=**
```

The latter (CATDSN=**) will select all cataloged data sets, so it should always be used with VOL= or STORGRP= to limit it to the volumes which will be cloned.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

CATALOG=

Specifies the name of a user catalog which FDRCLONE is to update. Only data sets in that user catalog which match the CATDSN and VOL/STORGRP operands will be updated.

Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter. CATDSN=** will search the master catalog plus all user catalogs with aliases assigned.

50.74 CONTINUED . . .

CATLIMITGDG= May be used to limit the cloning of GDGs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:

n will cause only the most recently created "n" generations to be cloned.

-n will cause only generation (-n) to be cloned.

Default is that all the generations of cloned GDGs will be selected unless a relative generation number is specified at the end of the filter, e.g.,
CATDSN=filter(-2).

COPY= Specifies the copy of the backup from which the restore of the cloned data sets is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utility ([See Section 60](#)). COPY= is ignored for CLONE TYPE=ACTIVATE and DECLONE.

If COPY=1 or 2 and ABR finds that one of the backup tapes is not cataloged under the copy specified, ABR will check to see if the other copy was created. If cataloged, ABR will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) ABR will automatically use the other copy (1 or 2) if it still exists.

Default is the value from the COPY statement or from the FDR Global Option table ([See Section 90](#)). For CLONE TYPE=DEFER, the option table in effect at the time the TYPE=DEFER job is run will provide the default.

GEN= Specifies the ABR generation from which the data sets named on this SELECT statement will be cloned; they will be restored from the full-volume backup (CYCLE=00) in that generation. This is used when you want to clone from other than the most recently cataloged generation. If the SELECT statement selects data sets from multiple disk volumes (e.g., multiple volumes in the VOL= operand), all of them will be cloned from that generation of the volume to which they were cataloged.

The default is:

for **CLONE TYPE=INPLACE**, FDRCLONE will locate the most recently cataloged ABR full-volume backup (CYCLE=00 in the most recent generation) for each volume in each cataloged data set selected.

for **CLONE TYPE=DEFER**, FDRCLONE will locate the most recent backup of every selected data set, using the information recorded on the volume itself (in the ABR Model DSCB and the F1 DSCB of the data set itself). This backup may be on a full-volume backup (CYCLE=00) or an incremental backup (CYCLE>00).

Note: Unless the site has taken steps to keep the ABR generation number of related volumes in sync, it is likely that different volumes will have different generation numbers at any given point-in-time. If so, the only circumvention is to code multiple SELECT statements, e.g.,

```
SELECT CATDSN=**, VOL=PROD01, GEN=75
SELECT CATDSN=**, VOL=PROD02, GEN=112
```

See the "Performance Note" above to understand the performance implications of this.

50.74 CONTINUED . . .

- PRTALIAS** Will display all of the alias names and user catalogs that were searched for this SELECT statement. This is effective only when the data set name mask you provided forces CATDSN= to start in the master catalog and search one or more alias catalogs. If the non-mask characters at the beginning of the mask are sufficient to cause CATDSN= to begin its search in a user catalog, PRTALIAS is ignored.
- STORGRP=** Specifies the name of an SMS storage group, up to 8 characters. You may specify multiple storage group names by enclosing them in parenthesis, e.g., STORGRP=(DB,PROD). It can be used only on systems with SMS (System Managed Storage) active. It is equivalent to VOL= for all the volumes defined to be part of that storage group.
- VOL=** Specifies a volume serial number (e.g., VOL=ABC123, up to 6 characters) or a volume serial prefix followed by an asterisk (e.g., VOL=(ABC*)) to which this statement applies. You may specify multiple volume serials or prefixes by enclosing them in parenthesis, e.g., VOL=(ABC123,XYZ*). If necessary, you can break after any comma in the list and continue the list on another line.
- STORGRP= and VOL= are mutually exclusive. If neither is specified, FDRCLONE will clone all data sets which match the CATDSN= operand regardless of their volume serial.
- If STORGRP= or VOL= is specified with CATDSN=**, all data sets cataloged to those volumes will be marked for cloning. If CATDSN=mask is specified, only data sets which match the mask and are cataloged to those volumes are marked.

50.75 FDRCLONE EXAMPLES

**CLONE IN-
PLACE BY
VOLUME**

This will mark all cataloged data sets which are cataloged to volumes starting with PROD or TSO for cloning. For best performance, list **all** volume serials on the same SELECT statement (up to 255). This CLONE TYPE=INPLACE job will be executed on the clone target system after all catalogs (including the ABR catalog) have been restored. The TAPE1 DD will contain a copy of the modified catalog entries *before* modification; the high-level index of FDRABR will insure that it is cataloged in the ABR catalog, but you can use any naming convention you like; if it is a GDG, a GDG base must be created. COPY=2 indicated that the cloned data sets will be restored from COPY2 of the backups (COPY2 is usually the offsite copy).

Note: CLONE TYPE=INPLACE will only mark data sets for cloning from the most recent full-volume backup of the volumes to which they are cataloged. Since the data sets do not exist on disk at the clone target system, it cannot determine which incremental backups contain backups of specific data sets.

```
//FDRCLONE      EXEC    PGM=FDRCLONE,REGION=0M
//SYSPRINT      DD      SYSOUT=*
//TAPE1         DD      DSN=FDRABR.FDRCLONE(+1),UNIT=CART,
//              DISP=(,CATLG),LABEL=EXPDT=99000
CLONE  TYPE=INPLACE, COPY=2
SELECT CATDSN=**,VOL=(PROD*,TSO*)
SELECT CATDSN=**,VOL=(PROD*,TSO*)
```

```
FDR001  FDR CLONE PROGRAM          - FDRCLONE  VER. 5.3/30P  - INNOVATION DATA PROCESSING
FDR303  CARD IMAGE  --  CLONE TYPE=INPLACE
FDR303  CARD IMAGE  --  S CATDSN=**,VOL=(PROD*,TSO*)
FDR346  FDR CLONED VOL=PROD01 DSN=ICF091.ESDS.DEFAULT.DATA          <-- GEN=0245 CYCLE=00
FDR346  FDR CLONED CLUSTER=ICF091.ESDS.DEFAULT.CLUS
FDR346  FDR CLONED VOL=PROD01 DSN=PAYROLL.A0007.DATA          <-- GEN=0245 CYCLE=00
FDR346  FDR CLONED VOL=PROD02 DSN=PAYROLL.A0007.INDEX          <-- GEN=0245 CYCLE=00
FDR346  FDR CLONED VOL=PROD87 DSN=PAYROLL.AIX.A0007.DATA          <-- GEN=0122 CYCLE=00
FDR346  FDR CLONED VOL=PROD87 DSN=PAYROLL.AIX.A0007.INDX          <-- GEN=0122 CYCLE=00
FDR346  FDR CLONED CLUSTER=PAYROLL.A0007.CLUS
FDR346  FDR CLONED VOL=TSO100 DSN=BAB2.ISPF.PROFILE          <-- GEN=0035 CYCLE=00
FDR346  FDR CLONED VOL=TSO110 DSN=ISPF.CLIST          <-- GEN=0035 CYCLE=00
FDR346  FDR CLONED VOL=PROD12 DSN=ACCOUNT.MASTER.G0003V00          <-- GEN=0262 CYCLE=00
FDR346  FDR CLONED VOL=PROD12 DSN=ACCOUNT.MASTER.G0004V00          <-- GEN=0262 CYCLE=00
FDR346  FDR CLONED VOL=PROD15 DSN=ACCOUNT.MASTER2.G0007V00          <-- GEN=0017 CYCLE=00
FDR346  FDRCLONE CLONED          7 FILES/CLUSTERS
```

This is an example of the output of this FDRCLONE job

For non-VSAM data sets, one FDR346 message is printed for each data set. However, if the data set is cataloged to multiple volumes, and more than one volume in the catalog is being cloned, the FDR346 is repeated for each volume in the data set.

For VSAM clusters, every component of the cluster (including alternate index components) which is cataloged to a cloned volume will be listed (multiple times for multi-volume components). Following that the cluster name is listed.

GEN= identifies the generation number of the ABR full-volume backup from which the data set or component will be restored.

The final FDR346 message totals the number of VSAM clusters (not components) and non-VSAM data sets which were cloned. Of course, in a real FDRCLONE execution, this number will be much larger.

50.75 CONTINUED . . .

**CLONE
DEFERRED BY
STORAGE
GROUP**

This will mark all cataloged data sets which are cataloged to volumes defined in the SMS storage group DB1 and FINANCE for cloning. This CLONE TYPE=DEFER job is run on your production system after ABR backups have been run. It does not update the current system. It will write a "before and after" image of every selected catalog record to the data file defined by TAPE1; this data file will be used on the clone target system to update the catalogs to mark data sets for cloning. Both full-volume and incremental ABR backups are supported. The high-level index of FDRABR will insure that it is cataloged in the ABR catalog, but you can use any naming convention you like; if it is a GDG, a GDG base must be created. COPY=2 indicates that the cloned data sets will be restored on the clone target system from COPY2 of the backups (usually the offsite copy).

```
//FDRCLONE EXEC PGM=FDRCLONE,REGION=0M
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DSN=FDRABR.FDRCLONE(+1),UNIT=CART,
// DISP=(,CATLG),LABEL=EXPDT=99000
CLONE TYPE=DEFER,COPY=2
SELECT CATDSN=**,STORGRP=(DB1,FINANCE)
```

**CLONE
ACTIVATE BY
STORAGE
GROUP**

This is executed on the clone target system, after all catalogs have been restored. It reads the data file produced by the previous example and updates the catalogs to mark the selected data sets as "cloned". Note that the SELECT statements in this step must match those in the TYPE=DEFER step which produced the data file; however, you can select a subset of the original data sets if you only want to clone that subset.

Note: if you are cloning by SMS storage group name, the SMS storage group on the clone target system must be defined to contain all of the original volumes in that group on the original system. Otherwise, FDRCLONE will not select the proper entries unless you use VOL= to specify the original volumes.

```
//FDRCLONE EXEC PGM=FDRCLONE,REGION=0M
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DSN=FDRABR.FDRCLONE(0),DISP=OLD
// DISP=(,CATLG),LABEL=EXPDT=99000
CLONE TYPE=ACTIVATE
SELECT CATDSN=**,STORGRP=(DB1,FINANCE)
```

**CLONE
IN-PLACE BY
DATA SET**

This will mark all cataloged data sets which match the data set name selection masks for cloning, regardless of the volumes to which they are cataloged. FDRCLONE will restore data sets from the backup copy (1 or 2) indicated by the BKPCOPY option in the FDR Global Option Table.

```
//FDRCLONE EXEC PGM=FDRCLONE,REGION=0M
//SYSPRINT DD SYSOUT=*
CLONE TYPE=INPLACE
SELECT CATDSN=PAYROLL.**
SELECT CATDSN=ENG**
```

**CLONE
IN-PLACE BY
DATA SET AND
VOLUME**

This will mark all cataloged data sets that are cataloged to the specified volumes and match the data set name selection mask for cloning. FDRCLONE will restore data sets from the backup copy (1 or 2) indicated by the BKPCOPY option in the FDR Global Option Table.

```
//FDRCLONE EXEC PGM=FDRCLONE,REGION=0M
//SYSPRINT DD SYSOUT=*
CLONE TYPE=INPLACE
SELECT CATDSN=PAYROLL.** ,VOL=(PAY*,PRODOF)
```

50.75 CONTINUED . . .

**APPLICATION
BACKUP AND
RESTORE OF
CATALOGS**

FDRCLONE requires that you backup on the production system all catalogs that contain data sets to be cloned and the ABR catalog, and restore them on the clone target system. The following is an example of using FDRAPPL (application backup) to do so. It assumes that all of your catalogs can be identified by name, such as that they all start with "CATALOG."; if this is not true, you may have to list them on separate SELECT statements. More details on FDRAPPL are found in [Section 52](#). You can also backup and restore the catalogs with FDRDSF, but this requires that you identify all of the volumes containing the catalogs and setup a separate backup for each; FDRAPPL automates this process.

This BACKUP step requires that the GDG bases for
TECH.DR.ACFBKP1 and
TECH.DR.ACFBKP2

be previously built with an appropriate number of generations. For disaster/recovery, you should send one and perhaps both of the ACFBKP tapes offsite for use at the DR site.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSPRINT1   DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=TECH.DR.CATALOGS,DISP=(NEW,DELETE),
//              UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)
//TAPE1       DD     DSN=TECH.APPL1,UNIT=CART,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=EXPDT=99000
//TAPE11      DD     DSN=TECH.APPL2,UNIT=CART,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=EXPDT=99000
//SYSIN       DD     *
              DUMP TYPE=APPL,ARCB1DSN=TECH.DR.ACFBKP1(+1),
              ARCB2DSN=TECH.DR.ACFBKP2(+1)
              SELECT CATDSN=CATALOG.**
```

The RESTAF step will restore only the application control file (ACF) created by the FDRAPPL backup. Since the catalogs have not yet been restored, you must specify the full data set name, tape volume serial and file sequence number of the ACF backup file; these can be found in the output of the BACKUP step.

Step RESTAPPL will restore all of the catalogs.

Note that the ACF and the catalogs will be restored to their original volume serials unless you have provide a RESTORE Allocation List or added NVOL operands to the SELECT statements. If the catalogs do not already exist on the target volumes, you must be sure that those catalog names do not exist in the master catalog of the clone target system; use IDCAMS EXPORT DISCONNECT to delete them if necessary.

This example restores from the COPY1 backup (TAPE1 in the BACKUP example above). If you are restoring from COPY2 (TAPE11), change ACFBKP1 to ACFBKP2 and add ",COPY=2" to the RESTORE TYPE=APPL statement.

```
//RESTAF      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=TECH.DR.ACFBKP1.G0175V00,UNIT=CART,
//              VOL=SER=024387,LABEL=10,DISP=(OLD,PASS)
//SYSIN       DD     *
              RESTORE TYPE=DSF,RECAT
              SELECT DSN=TECH.DR.CATALOGS
//RESTAPPL    EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=TECH.DR.CATALOGS,DISP=SHR
//TAPE1       DD     DSN=TECH.DR.ACFBKP1.G0175V00,DISP=(OLD,KEEP)
//SYSIN       DD     *
              RESTORE TYPE=APPL
              SELECT ALLDSN
```

50.80 FDRDRP

FDRDRP FDRCLONE also includes FDRDRP (Disaster Recovery Program), a utility for optimizing full-volume recovery from ABR volume backups (full-volume and incremental backups). It is intended for use at disaster/recovery sites. This can be used for volumes where full-volume recovery is more appropriate, such as system volumes and high-priority production volumes which must be recovered quickly. Cloning can still be used in conjunction with FDRDRP, for volumes and data sets which are not so urgent and those which must be spread out over different volumes.

FDRDRP recovers the complete image of each disk volume like a jigsaw puzzle. Different pieces of the backup of each volume (the ABR incremental and full-volume backups) may be read at various times, with restore activity for other volumes interspersed, but FDRDRP will assemble the backups in their proper places on disk to reconstruct the image of the original disk volume. Even though the backup tapes contain backups for many disk volumes, the tapes are mounted a minimum number of times.

**ABR VOLUME
RECOVERY**

ABR full-volume recovery starts by reading the most recent ABR cycle (incremental backup) for a given disk volume, then reads the next most recent, etc., until it reads the most recent full-volume backup (cycle 00) of that disk and completes the recovery. Since these backups are probably on separate tapes, the restore usually mounts and unloads a backup tape for each cycle to be read. [See Section 50.01](#) for more information on volume backups, generations, cycles, and full-volume recovery

Although a ABR full-volume recovery step may request the restore of many disk volumes, a normal ABR full-volume recovery restores one disk volume at a time. Since a given backup tape may contain the backups of many disk volumes, the same tapes may be unloaded and remounted over and over, taking considerable time and overwhelming operators and automated tape libraries.

Even if ABR recoveries are run in separate jobs, allowing the restores to be run in parallel, the same tape mounts will occur and the restore jobs may contend for the same backup tape volumes, so the situation is no better. Multiple restore jobs are useful only when each uses a different set of input tapes.

This is particularly a problem on modern high-capacity tapes such as the IBM Magstar and STK Redwood and 9840, since they can hold the backups of many disk volumes. High-capacity tapes use fewer tape volumes to hold your backups, but those volumes must be mounted over and over.

But even on lower-capacity tapes, such as 3490E, tapes may need to be remounted many times during the restore of the requested disk volumes.

50.80 CONTINUED . . .

**THE FDRDRP
SOLUTION**

FDRDRP processes multiple full-volume recovery tasks in parallel. It manages usage of the backup tapes required for those restores, so that each backup tape is mounted a minimum number of times, **usually one mount per tape volume**. This will greatly reduce the elapsed time required to recover the volumes and eliminate most extra tape mounts required by ABR.

The nature of tapes is that only one task can use a given tape volume at a time. FDRDRP manages the use of tapes by this process (slightly simplified):

- a recovery subtask is started for each volume specified by a SELECT statement in the FDRDRP input. However, the disk volumes are sorted by the tape volume serial and file sequence number required for the first backup each will read, so that the subtasks will read the backup files on a tape in physical order with minimal positioning
- a restore subtask dynamically allocates and mounts the first backup tape it requires
- if additional restore subtasks require other backup files on the same tape, they will wait on the owning subtask to release it. As a subtask finishes with a tape volume it passes the tape to a waiting subtask without dismounting the tape. If no other subtask is waiting for that tape, the tape is no longer needed and is deallocated and dismounted (unloaded). If more than 100 disk volumes are requested, FDRDRP will restore them in groups of 100 volumes at a time.
- if the subtask that just passed a tape volume needs another tape volume, it will allocate and mount it (or wait if another subtask happens to be using that volume). If the restore is complete (after reading cycle 00, the full-volume backup), the subtask will terminate
- when a restore subtask requires a backup tape that is not currently in use by another subtask, it will allocate and mount the tape. However, the MAXTAPES= operand on the RESTORE statement should be used to limit the total number of tape drives, allowing you to specify the number of drives you are willing to devote to the FDRDRP restore step. If a task requires a backup tape not currently in use but MAXTAPES drives are already in use by this FDRDRP step, those tasks will wait until the count of active tape drives decreases
- once all restore subtasks have completed, the FDRDRP step will terminate
- if you are running multiple FDRDRP jobs, each restoring a different set of disk volumes, but two such jobs happen to require the same tape volume at the same time, FDRDRP will pass the volume from one job to the other without dismounting it.

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

50.80 CONTINUED . . .

**FDRDRP
OPERATION
EXAMPLES**

Here is a simple example to show how FDRDRP operates. Three disk volumes have been backed up, once by full-volume and twice by incremental. The offsite COPY 2 backup tapes look like:

Tape Volume:	333333	222222	111111
File 1:	FDRABR.VPROD01. C2012202	FDRABR.VPROD01. C2012201	FDRABR.VPROD01. C2012200
File 2:	FDRABR.VPROD02. C2001702	FDRABR.VPROD02. C2001701	FDRABR.VPROD02. C2001700
File 3:	FDRABR.VPROD03. C2000302	FDRABR.VPROD03. C2000301	FDRABR.VPROD03. C2000300

Each day's backups are all contained on one tape. The typical sequence of FDRDRP operation is:

Disk Volume:	PROD01	PROD02	PROD03
Time T:	Mount tape 333333, restore from file 1	Wait for volume 333333	Wait for volume 333333
T+1:	Mount tape 222222, restore from file 1	Restore from file 2 on 333333	
T+2:	Mount tape 111111, restore from file 1	Restore from file 2 on 222222	Restore from file 3 on 333333 and dismount
T+3:	(full-volume backup)	Wait for volume 111111	Restore from file 3 on 222222 and dismount
T+4:	Volume Restore completed.	Restore from file 2 on 111111	Wait for volume 111111
T+5:		(full-volume backup)	
T+6:		Volume Restore completed.	Restore from file 3 on 111111 and dismount
T+7:			(full-volume backup)
T+8:			Volume Restore completed.

An ABR full-volume recovery of these same 3 disks would mount each tape 3 times (9 mounts) and would take time to position to the required file. FDRDRP mounted each tape only once (3 mounts) and eliminated positioning delays, resulting in a typical elapsed time saving of over 80%.

The shaded blocks show where each task is waiting for a tape to become available. As of time T+2, 3 tape drives are in use and all 3 restore subtasks are actively restoring data (this can be limited by the MAXTAPES= operand). At times T+1 and T+3 two tapes are in use and two subtasks are restoring data.

Since the full-volume backups are usually much larger than the daily incrementals, it takes longer to read it, so the restores will eventually serialize on that tape volume.

50.80 CONTINUED . . .

However, full-volume backups frequently require more than one tape volume (especially when lower-capacity tapes such as 3480 and 3490E are used) and this actually improves FDRDRP efficiency. Here is an example showing restore from only the full-volume backups (no incrementals):

Tape Volume:	File 1	File 2	File 3
444444	FDRABR.VTSO001. C2011200	FDRABR.VTSO002. C2001400 starts	
555555		FDRABR.VTSO002. C2001400 concluded	FDRABR.VTSO003. C2000300 starts
666666			FDRABR.VTSO003. C2000300 ends

So the backups of TSO001 and TSO003 are entirely contained on one tape each, while TSO002 exists partially on both tapes. The typical sequence of FDRDRP operation in this example is:

Disk Volume:	TSO001	TSO002	TSO003
Time T:	Mount tape 444444, restore from file 1	Wait for volume 444444	Mount tape 555555, restore from file 3 and dismount
T+1:	Volume Restore completed.	Wait for volume 555555	
T+2:		Restore from file 2 on 444444 and dismount	Mount tape 666666, continue restore from file 3 and dismount
T+3:		Mount tape 555555, continue restore from file 2 and dismount	Volume Restore completed.
T+4:		Volume Restore completed.	

Because disk volumes TSO001 and TSO003 don't require any common tapes, they proceed in parallel. TSO002 has a common tape volume with both other disk volumes so it will wait for both restore tasks to release the tapes it needs. Tape 555555 will be mounted twice.

These examples are very simple. In a real restore, where each day's backups may be on a variety of tapes or on multi-volume tape sets, the sequence will be more complicated. There may even be conditions where FDRDRP must release a tape and remount it later. However, no matter how complex the restores, FDRDRP will maximize the number of concurrent restores (subject to MAXTAPES=) and minimize the number of tape mounts, greatly reducing restore elapsed time compared to ABR full-volume recoveries from incremental backups.

50.80 CONTINUED . . .

**FDRDRP
CONSIDER-
ATIONS**

Each FDRDRP restore subtask must read the backup tapes in the normal order used by ABR full-volume recovery, reading the most recent incremental backup first, then the next oldest, etc., until the full-volume backup is read.

FDRDRP works efficiently when all the backups for a given day for each disk volume being restored are on one tape volume or multi-volume set. This is the pattern shown in the example above. Each input tape is mounted only once (it may be necessary to mount a tape twice if a backup file crosses from one tape volume to another).

This is compatible with the way that most ABR users run their volume backups. A typical installation will run ABR volume backups daily, selecting all the volumes to be backed up. Scratch tapes are used for output and ABR automatically stacks backup files on tape, so the output tape (or multi-volume tape set) will contain all the backups created on that day. If you use multiple TAPEX DD statements in the ABR step or run multiple ABR backup jobs, multiple tapes or tape sets are created containing backups for a subset of your disk volumes but they still contain only backups created on that day. In most cases, FDRDRP will be able to restore all the volumes while mounting each backup tape only once.

FDRDRP will work efficiently even if you use the ABR LASTAPE feature to add backup files from multiple days onto an existing tape or tape set.

However, if your backups are not so neatly ordered, perhaps because you don't backup every disk every day or some backups failed, the restore is not so simple. For example, if the first (latest) cycle required by one subtask is on the same tape as the third cycle required by another subtask, that subtask may not be ready to read that tape when the first subtask is finished with it. In that case, the tape may be dismounted and remounted later. **This is unusual and will not occur for most FDRDRP users.** Dismounting and remounting may also occur if you are restoring a large number of disk volumes in one FDRDRP step.

Note: if you invoke concurrent ABR backups by multiple TAPEX DD statements in the ABR JCL, the choice of which disk volumes will go to which TAPEX DD is dynamic and may vary from day to day depending on the amount of data backed up from each volume. This will complicate the FDRDRP restore process and may slow it down and/or require more tape mounts.

FDRDRP will execute more efficiently if you specify MAXTAPES= equal to the number of tape drives you have available at the disaster site (or the number of drives you are willing to allocate to this FDRDRP step). If it is set to a value larger than the number of drives, extra overhead is incurred while the FDRDRP subtasks contend for the available tapes.

50.80 CONTINUED . . .

**FDRDRP
CONSIDER-
ATIONS**

MAXFILE=1 Note: Some customers run their full-volume ABR backups with MAXFILE=1 so that each disk volume starts its backup as file 1 on a fresh scratch tape. This improves ABR restore efficiency when restores were run in parallel ABR jobs. With FDRDRP, MAXFILE=1 is not recommended. You can use the default of MAXFILE=255 or even specify a larger value, in order to fill tape volumes to capacity, and still get a great deal of restore parallelism with FDRDRP.

WARNING: FDRDRP does many dynamic allocations for the tapes it needs. Dynamic allocation of tape will be affected if there is an unsatisfied allocation recovery message on the console, such as

```
IEF244I jobname RESTORE - UNABLE TO ALLOCATE 1 UNIT(S)
IEF877E jobname NEEDS 1 UNIT(S)
FOR RESTORE TAPE05
FOR VOLUME 001310
OFFLINE,
0393-0395
IEF878I END OF IEF877E FOR jobname RESTORE TAPE05
98 IEF238D jobname - REPLY DEVICE NAME, 'WAIT' OR 'CANCEL'.
```

Until the operator replies to this message, no further allocations of that type of tape can be satisfied, which will probably cause FDRDRP to eventually wait until the reply is made. This is true even if the allocation recovery is for another non-FDR job. If the operator replies WAIT, message IEF433D must also be satisfied, replying HOLD or NOHOLD.

FDRDRP may occasionally issue a console UNLOAD (U) command for tapes it has mounted. These UNLOAD commands may also appear in the joblog of the FDRDRP job. This is normal. If a restore task is done with a tape and the next restore task needing that tape is not yet ready to accept it within a short period of time, it is unloaded so that the tape drive is not tied up unnecessarily. An UNLOAD usually means that the tape will be remounted during another part of the restore process.

Unless your tape management and security databases have been recovered to a point after the creation of these ABR backups at your home site, you are likely to get tape management or security errors as ABR tries to open the tapes. Innovation recommends that you disable tape management and security checking during the FDRDRP restores. If you cannot disable tape management, specify "EXPDT=98000" on the RESTORE statement; for most tape management systems that will bypass tape management checking at OPEN time; note that the use of 98000 may be limited by your security system or by options in the tape management system itself.

**TESTING
FDRDRP**

A full-blown test of FDRDRP normally must be done at the disaster/recovery site, where you can restore all of the disk volumes required. But you will probably want to verify that FDRDRP works before you devote the resources to that full-blown test.

You can execute a limited test of FDRDRP at your home site (or on an LPAR) if you have a number of disk volumes, such as unused volumes or scratch (temporary) volumes, which you can dedicate as target volumes for the FDRDRP restore test. You use FDRDRP to restore a subset of your production volumes to those temporary target volumes. Although FDRDRP will work correctly with as little as 1 tape drive, it will be a more realistic test if you can dedicate several tape drives to the test.

An example is shown in [Section 50.83](#).

50.81 FDRDRP JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform restores from ABR Volume Backups with FDRDRP:

**STEPLIB or
JOB LIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRDRP resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRDRP), region requirement (REGION=), and optional PARM= operand. REGION=0M is recommended to get the largest below-the-line region available.

If a PARM field is specified, FDRDRP will use data specified as the first control statement, which must be a valid RESTORE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRDRP,PARM='RESTORE TYPE=DRP,MAXTAPES=3'
//FDR EXEC PGM=FDRDRP,PARM='RESTORE TYPE=DRP/ SELECT VOL=PROD01'
```

If FDRDRP is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**SYSPRINT DD
STATEMENT**

Specifies the output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

FDRDRP will also dynamically allocate a SYSOUT=* print data set for each disk volume being restored, using the DD name SYSPRTnn. Although SYSPRTnn will contain the complete messages associated with the restore of a disk volume, SYSPRINT will also contain many of those messages as well, grouped together by disk volume, to make it easier to locate the results for a specific volume and to spot error messages.

**FDRSUMM DD
STATEMENT**

(Optional) if present, FDRDRP will write one-line messages for each volume restored, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set.

**SYSUDUMP DD
STATEMENT**

Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for FDRDRP. Usually a DD * data set.

It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

50.82 FDRDRP RESTORE STATEMENT

```

RESTORE      TYPE=DRP                      ,MAXERR=nn
R
              ,CONFMESS=YES|NO             ,MAXTAPES=nn
              ,COPY=n                     ,OPERATOR
              ,CPYVOLID=NO|YES             ,SMSPROT=ALL|NONE
              ,EXPDT=yyddd|yyyddd         ,VOLRESET=YES|NO

```

**RESTORE
STATEMENT**

The RESTORE TYPE=DRP statement performs a full-volume recovery of one or more DASD volumes from ABR Volume Backups. Only one RESTORE statement is allowed per execution of FDRDRP.

WARNING: before you do a full-volume restore you should make sure that the target volume is offline to all systems other than the system where the restore is to be run. If you don't, the other systems may access the original VTOC of the restored volume and access or delete the wrong data.

OPERANDS

TYPE=DRP Required for a FDRDRP full volume restore. It is followed by one or more SELECT statements, each specifying a volume to be recovered.

CONFMESS= **YES** — before beginning the restore, FDR will request confirmation via a WTOR (FDRW01) message to which the MVS operator must reply.
NO — suppresses the WTOR and begins the restore immediately.
 Default: NO.

COPY= Specifies the copy of the backup from which the restore is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utility ([See Section 60](#)).

If COPY=1 or 2 and ABR finds that one of the backup tapes is not cataloged under the copy specified, ABR will check to see if the other copy was created. If cataloged, ABR will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) ABR will automatically use the other copy (1 or 2) if it still exists.

Default is COPY=1 unless overridden in the FDR Global Option table ([See Section 90](#)).

50.82 CONTINUED . . .

- CPYVOLID=** Specifies whether the volume serial number of the disk that was backed up will be restored, if the existing volser of the output disk is different (if the serials are the same, CPYVOLID is ignored).
- YES**— volume serial number of the output volume will be replaced with the original volume serial number of the disk which was dumped. If another online volume has the same serial, the restored volume will be placed offline at the end of the restore.
- NO**— the volume serial number of the output volume will be retained. See the VOLRESET= operand below.
- Default: NO — unless the volume being restored was SMS-managed, when YES is forced.
- Note:** although full-volume ABR restore does not catalog data sets, any data sets which were cataloged to the original volume are automatically cataloged to the new volume when restoring with CPYVOLID=YES, assuming that you restore the system catalogs to the same point-in-time. If you use CPYVOLID=NO and do not later relabel the volume, data sets may need to be manually recataloged.
- VSAM/SMS WARNING:** *See the notes under VOLRESET=NO below. CPYVOLID=YES is recommended for any volume containing a VVDS.*
- EXPDT=** Specifies an expiration date which will be passed to OPEN when each backup file is opened. Since these are input tapes, the expiration date will probably be ignored by your tape management system except for certain special dates. The most common use will be EXPDT=98000, which is accepted by most tape management systems and means "this tape is not in the tape management data base". You might need to use EXPDT=98000 (or whatever your TMS supports) when the tape management data base has not been restored to a point after the ABR backups were taken, so that it doesn't reflect the backup tapes that FDRDRP needs to read. In this case, it might be easier to disable tape management until the data base can be made current.
- The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).
- MAXERR=** Specifies the number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP.
- Default: 20 errors.

50.82 CONTINUED . . .

CPYVOLID= Specifies whether the volume serial number of the disk that was backed up will be restored, if the existing volser of the output disk is different (if the serials are the same, CPYVOLID is ignored).

YES— volume serial number of the output volume will be replaced with the original volume serial number of the disk which was dumped. If another online volume has the same serial, the restored volume will be placed offline at the end of the restore.

NO— the volume serial number of the output volume will be retained. See the VOLRESET= operand below.

Default: NO — unless the volume being restored was SMS-managed, when YES is forced.

Note: although full-volume ABR restore does not catalog data sets, any data sets which were cataloged to the original volume are automatically cataloged to the new volume when restoring with CPYVOLID=YES, assuming that you restore the system catalogs to the same point-in-time. If you use CPYVOLID=NO and do not later relabel the volume, data sets may need to be manually recataloged.

VSAM/SMS WARNING: See the notes under VOLRESET=NO below.
CPYVOLID=YES is recommended for any volume containing a VVDS.

EXPDT= Specifies an expiration date which will be passed to OPEN when each backup file is opened. Since these are input tapes, the expiration date will probably be ignored by your tape management system except for certain special dates. The most common use will be EXPDT=98000, which is accepted by most tape management systems and means "this tape is not in the tape management data base". You might need to use EXPDT=98000 (or whatever your TMS supports) when the tape management data base has not been restored to a point after the ABR backups were taken, so that it doesn't reflect the backup tapes that FDRDRP needs to read. In this case, it might be easier to disable tape management until the data base can be made current.

The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

MAXERR= Specifies the number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP.

Default: 20 errors.

50.82 CONTINUED . . .

MAXTAPES= Specifies the maximum number of tape drives (1 to 99) which FDRDRP will use to restore the selected disk volumes. You can use this to insure that other tape drives remain free for other recovery operations.

The default is 10.

Note: If FDRDRP needs another tape drive (MAXTAPES not reached) and no additional tapes of the proper type are online, it will issue console message FDRW10 asking you to vary an additional drive online. If you do so, FDRDRP will recognize it and begin using it after a few seconds; if you ignore the message, FDRDRP will only use tape drives which are already online. You may want to insure that MAXTAPES tape drives are online before you start the FDRDRP job to avoid this interaction. If the disaster system has fewer than 10 tape drives of the proper type, you should specify MAXTAPES=n (the actual number available) to avoid extra overhead in FDRDRP.

OPERATOR Specifies that, before the RESTORE operation began, an operator message will be issued for each tape necessary to complete the restore. This option gives the operator the ability to pre-pull required tapes or bypass individual cycles for which the tapes may not be available at this time. However, if some cycles are bypassed, the restore may correctly restore the latest version of some data sets.

Note: since the restore subtasks for all the selected disk volumes will begin at the same time (up to 100 at a time), the FDRW25 messages generated by OPERATOR will appear for all disk volumes in a very short period, which may be very confusing and will require a lot of replies. It is not recommended. However, if some off-site tapes are unavailable, this may be the only way to bypass the missing tapes.

SMSPROT= **ALL** — enforces several rules when SMS-managed volumes are involved: Backups of SMS-managed volumes can only be restored to SMS-managed volumes, and non-SMS volumes only to non-SMS volumes. CPYVOLID=YES is forced when an SMS-managed volume is restored.

NONE — allows the restore of SMS-managed volumes to non-SMS volumes, and vice versa. Also allows the restore of SMS volumes to new volumes if CPYVOLID=NO is specified.

WARNING: SMSPROT=NONE will usually be used at a disaster recovery site where a relPL of MVS will be done after the restores, to place all volumes in the proper SMS status. See Section 70 for more details on restoring and moving SMS-managed volumes.

Default: NONE.

50.82 CONTINUED . . .

VOLRESET= When CPYVOLID=NO is specified or defaulted, and the volume serial of the output disk is different from that of the original disk on the backup data set, the backup is restored but the volume serial of the output volume is retained. VOLRESET= controls additional processing relating to this change in volume serials. VOLRESET is ignored if the volume serial of the output disk is not being changed.

YES — the volume serials that are part of the data set names of the VTOC Index ("SYS1.VTOCIX.volser") and the ABR Model DSCB (usually "FDRABR.Vvolser") are checked to see if they match the input volume serial (the volume backed up). If so, they are renamed to match the volser of the output disk. Also, the DSCB field DS1DSSN (data set serial number, usually the volser of the first or only volume of the data set) for every data set on the volume will be changed to the new volume serial if the existing value matched the original volume serial. In the ABR Model DSCB, the current cycle number is set to the cycle number of the first incremental read during the restore, so that the next incremental will produce the next cycle in this generation (a restore using TAPEDD= will set the cycle number to 63 to force the next backup to be a full-volume backup). VOLRESET=YES should be used with CPYVOLID=NO when you intend to permanently retain the volume serial of the output disk. Note that data sets on the volume will not be recataloged to the new volume serial.

NO — do not rename the VTOCIX and ABR Model DSCB, do not change DS1DSSN fields, and do not reset the cycle number in the ABR Model DSCB. VOLRESET=NO should be used with CPYVOLID=NO if you plan to eventually relabel the disk back to the original volume serial.

VSAM/SMS WARNING: FDRDRP will not rename the VVDS, since the VVDS and the catalogs contain self-defining records that would also need resetting. If a volume containing VSAM clusters or SMS-managed data sets is restored with CPYVOLID=NO, the data sets will be inaccessible unless the volume is relabeled to the original serial number. CPYVOLID=YES is recommended for any volume containing a VVDS, but if it is necessary to restore the volume under a temporary volser, use CPYVOLID=NO and VOLRESET=NO and relabel the volume with ICKDSF later.

Default: YES but it is ignored if CPYVOLID=YES is also specified.

50.83 FDRDRP SELECT STATEMENT

```

SELECT          VOL=vvvvvv          ,GEN=nnnn|CURRENT
S
                ,COPY=n              ,NVOL=vvvvvv
                ,CYCLE=nn

```

**SELECT
VOLUME
STATEMENT**

One or more SELECT Volume statement must follow a RESTORE TYPE=DRP statement to specify the volumes to be recovered, one per statement. Multiple SELECTs will cause those volumes to be restored in parallel; the restores will be initiated in an order sorted by the first tape volume serial each restore requires, and by the file sequence on that tape, so that tapes will be read in physical sequence.

You may include up to 1000 SELECT statements (1000 disk volumes) in a given FDRDRP step. If you have more than 100 disk volumes to restore, FDRDRP will do them in groups of 100 volumes. If you require more than 1000 volumes in a FDRDRP step, contact Innovation for assistance.

OPERANDS

VOL= Specifies the disk volume serial number of the original disk to be restored. It must be present. If NVOL= is not specified, the backup will be restored to an online disk volume which currently has the original serial (either the original disk or a new volume which has been initialized with that serial).

COPY= Specifies the copy of the backup from which the restore is to be done; "n" is a digit from 1 to 9. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created at backup time. Copies 2 through 9 can be created by the FDRTCOPY or FDRTSEL utility ([See Section 60](#)).

If COPY=1 or 2 and FDRDRP finds that one of the backup tapes is not cataloged under the copy specified, FDRDRP will check to see if the other copy was created. If cataloged, ABR will use the other copy. So, if the specified copy has expired (and been uncataloged by a tape management system) FDRDRP will automatically use the other copy (1 or 2) if it still exists.

Default is the value from the RESTORE statement or from the FDR Global Option table ([See Section 90](#)).

CYCLE= Specifies the cycle (incremental backup) number which FDRDRP is to read first while restoring this volume. This option is used if you do not wish to start the restore with the most recent cycle created in the current generation or if you specified GEN=gggg to restore from an older generation (CYCLE= is required if GEN= is given). FDRDRP will read this backup file first and continue backwards through the incrementals until the full volume (TYPE=FDR) backup is encountered. CYCLE=0 will only restore the full volume backup.

Default: see GEN= below.

50.83 CONTINUED . . .

GEN=

Specifies the generation number which FDRDRP is to read while restoring this volume. May be used if you wish to restore from a specific generation other than the most current. If GEN= is specified, CYCLE= must also be specified.

If GEN=CURRENT is specified, FDRDRP will search for the most recently created generation and cycle cataloged for this volume in the ABR catalog; the ABR Model DSCB is not used and original volume need not be online. If GEN=CURRENT,CYCLE=nn is specified, then the specified cycle in the generation from the catalog will be used (CYCLE=00 can be used to restore only the full-volume backup from the current generation).

By default, FDRDRP will read the ABR Model DSCB from the volume specified by VOL= to obtain the most recently created generation and cycle. If the volume is offline, if the model does not exist or has been destroyed, or if the model indicates that no backups exist, then it will default to GEN=CURRENT and get the GEN/CYCLE from the ABR catalog; since FDRDRP is intended to be run at a disaster site, this is the usual case. So, if you want to restore from the latest incremental backup in the current generation, GEN= should be omitted.

NVOL=

Specifies the current volume serial of the online disk volume to which this backup will be restored. This might be used at a disaster/recovery site where the site personnel pre-initialize all of the disks to known volume serials; you can do the FDRDRP volume recoveries without having to relabel all of their volumes to your volserials. If CPYVOLID=YES was specified or defaulted on the RESTORE statement, the output volume will be relabeled to the original volume serial at the end of the restore.

Default: the backup will be restored to an online disk labeled with the original volume serial.

50.84 FDRDRP RESTORE EXAMPLES

**FULL-VOLUME
RESTORE
FROM
CURRENT
BACKUP**

The 10 volumes indicated on the SELECT statements will be restored in parallel. Since the volumes at the disaster site are pre-initialized to a known volser, each SELECT specifies VOL= for the volume to be restored and NVOL= for the target volume to restore to. FDRDRP will locate the most-recently created COPY 2 backup for each disk volume in the ABR catalog, which must have been previously restored.

Each target volume will be relabeled to the serial of the restored disk. No more than 3 tape drives will be used by these restores. FDRDRP will mount the backup tapes the minimum number of times necessary to restore data for all disks. All required disks and tapes will be dynamically allocated, and FDRDRP will also dynamically allocate a print data set (DDname SYSPRTnn) to SYSOUT=* to contain messages from each restore task.

```
//RESTFULL EXEC PGM=FDRDRP,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
RESTORE TYPE=DRP,MAXTAPES=3,CPYVOLID=YES,COPY=2
SELECT VOL=PROD01,NVOL=DR0170
SELECT VOL=PROD02,NVOL=DR0171
SELECT VOL=PROD03,NVOL=DR0172
SELECT VOL=PROD04,NVOL=DR0173
SELECT VOL=PROD05,NVOL=DR0174
SELECT VOL=PROD06,NVOL=DR0175
SELECT VOL=PROD07,NVOL=DR0176
SELECT VOL=PROD08,NVOL=DR0177
SELECT VOL=PROD09,NVOL=DR0178
SELECT VOL=PROD10,NVOL=DR0179
```

**RESTORE OF
FULL-VOLUME
BACKUP ONLY**

A number of volumes will be restored in parallel from their most recent full-volume backup only. On each SELECT statement, GEN=CURRENT and CYCLE=00 selects the last full-volume backup.

Each target volume must be prelabeled with the serial of the restored disk. Up to 20 tape drives will be used by these restores if that many drives of the proper type are available. FDRDRP will mount the backup tapes the minimum number of times necessary to restore data for all disks. All required disks and tapes will be dynamically allocated, and FDRDRP will also dynamically allocate a print data set (DDname SYSPRTnn) to SYSOUT=* to contain messages from each restore task.

```
//RESTORE EXEC PGM=FDRDRP,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
RESTORE TYPE=DRP,COPY=2,MAXTAPES=20
SELECT VOL=TS0001,GEN=CURRENT,CYCLE=00
SELECT VOL=TS0005,GEN=CURRENT,CYCLE=00
.
.
.
SELECT VOL=TS0112,GEN=CURRENT,CYCLE=00
```

50.84 CONTINUED . . .

TEST FDRDRP Run a test of FDRDRP at your home site using unused or scratch volumes as targets. The 7 volumes indicated on the SELECT statements will be restored in parallel, using up to 3 tape drives. Each volume will be restored to an unused volume which has been reserved for this test. CONFMESS=YES is specified so that the operator can confirm that only the unused volumes are being modified. CPYVOLID=NO is specified so that the volume serial of the unused volumes will not be modified. FDRDRP will locate the most-recently created COPY 1 backup (incremental or full-volume) for each disk volume in the ABR catalog.

```
//TESTDRP      EXEC   PGM=FDRDRP,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
RESTORE        TYPE=DRP,CONFMESS=YES,CPYVOLID=NO,MAXTAPES=3
SELECT         VOL=PROD01,NVOL=SCR123
SELECT         VOL=PROD02,NVOL=SCR167
SELECT         VOL=PROD03,NVOL=SCR012
SELECT         VOL=PROD04,NVOL=TEMP52
SELECT         VOL=PROD05,NVOL=SCR001
SELECT         VOL=PROD06,NVOL=TEMP17
SELECT         VOL=PROD07,NVOL=NEW111
```

51.01 FDRABR ARCHIVING AND SUPERSCRATCH

ABR Program FDRABR (Automatic Backup Restore) automates the execution of FDR full-volume and data set backups and restores for the purposes of:

- **data availability** - creating backup copies of DASD data sets to protect against physical loss or logical damage.
- **space management** - identifying data sets which do not need to be on DASD (usually based on the last date they were used) and moving them to backups (called "archiving"), from which they can be automatically recalled if needed. Data sets which will never be needed again can be scratched without creating a backup.
- **disaster recovery** - creating backups from which all or part of your DASD data can be quickly recreated at a disaster recovery site.
- **tape mount management (TMM)** - maximizing tape volume usage by staging small tape data sets to a disk buffer, then moving many of them to a tape volume.

These objectives are accomplished by the use of several ABR functions:

- This section describes ABR ARCHIVE (DUMP TYPE=ARC) and SUPERSCRATCH (DUMP TYPE=SCR) which are used for "space management" and "tape mount management".
- ABR Volume Backups (DUMP TYPE=FDR/ABR/AUTO/DSF) are found in [Section 50](#).
- ABR Application Backup (DUMP TYPE=APPL) is in [Section 52](#).

ABR BACKUPS All the backup functions of ABR share these common characteristics:

- the backups are in standard full-volume (FDR) or data set (DSF) backup format, described in [Section 02.02](#). If necessary, data can be restored from ABR backups with FDR or FDRDSF, but ABR automates the restore process.
- although ABR backups require a separate backup data set for each DASD volume processed, ABR will automatically stack multiple backup data sets on tape, creating multi-file tapes, to make best use of today's high-capacity tape volumes (such as IBM Magstar and StorageTek Redwood and 9840). If necessary, multiple output tape volumes are used. No special JCL is required since ABR will handle the file creation internally.
- ABR backups to disk are also supported. ABR will automatically allocate the backup data sets on the backup volumes you specify. Disk backups are sometimes used with Archiving to provide fast recall of archived data, although customers with Automated Tape Libraries can get the same results with tapes in the ATL.
- the output devices (tape and/or disk) are specified in the ABR batch JCL. However, you only need to identify the output device; ABR will automatically name and create every backup data set and catalog it if required.
- You can specify in the JCL that two copies of each backup are to be created, even though the input disk is read only once. When Archiving this is often used to create one short-term copy on disk for fast recall and one on tape for long-term retention, but it can also be used to create a copy for offsite storage. In ABR these are known as COPY 1 and COPY 2 ([see Section 51.04](#) for details). The ABR utilities FDRTCOPY and FDRTSEL ([Section 60](#)) can be used to create COPY 2. Since an archived data set exists **only** in the backup files, 2 copies are recommended for Archive Backups.
- If multiple output devices are specified in the JCL, ABR will automatically use internal subtasking to process more than one DASD volume concurrently.

51.01 CONTINUED . . .

**ABR ARCHIVE
BACKUP**

Although the price per megabyte of mainframe disk storage has dropped considerably, the demand for disk storage has increased in most installations, so disk storage remains a significant part of the Data Center budget. Disk space must be used efficiently in order to justify its cost.

However, your disk volumes may contain data sets which no longer need to be on disk. In some cases, they are orphan data sets, created by some job or TSO user and never used again. Sometimes, they may be needed again in 6 months or a year, but don't need to be on disk in the meantime. Sometimes data sets simply don't meet installation standards, such as uncataloged data set.

ABR Archive Backup is used for space management, by moving data sets which are not currently required to be online from expensive disk to less expensive medium such as tape or highly-compressed disk. The data sets are moved by backing them up with an FDR data set backup, then deleting them from the disk. They can be recalled to disk whenever they are needed.

The DASD volumes to be processed by Archive must contain an ABR Model DSCB. [See Section 50.01](#) "ABR Volume Initialization" for details on creating and maintaining the ABR Model DSCB. Archive does not record anything in the ABR model, but there is a flag enabling or disabling each volume for archiving; you can protect whole volumes from archiving by disabling the flag, or, if you prefer, you can enable the flag only when you are about to run archiving on a given volume.

The data sets moved from DASD by ABR Archive are recorded in an Archive Control File (ACF). The ACF is a DSORG=DA (direct access) data set which can record several hundred such archived data sets in a single track. There is usually one common ACF for all archived data sets. For each archived data set, the ACF records the ABR backup file which contains the data set (device type, volume serials and file number), the expiration of that backup, and some basic descriptive information about the data set (DSORG, size, etc.). If two copies of the backup were made, both copies are recorded in the ACF. Since the ACF is cumulative, the most recent entries are recorded at its end.

You can specify various criteria for selecting the data sets to be archived, such as "not cataloged" or "not used in the last 15 days". Different criteria can be specified for various data sets (by volume, by data set name mask, or a combination). But the most commonly used criteria is days since last usage, since this automatically identifies data sets which have not been used recently and are probably candidates for archive. For SMS-managed data sets, you can optionally select data sets based on the attributes of the SMS management class assigned to each data set.

ABR Archive is often used to put one backup copy on disk and a second on tape. The disk backups will be in a highly compressed format, taking much less room than the original, yet they can be recalled quickly since no tape mount is required. Most installations expire the disk backup after a short time (15-30) days and either delete it or move it to tape. If you have an Automated Tape Library (ATL), restore from tape backups is almost as fast, so disk output is probably not needed.

**FDR
INSTANTBACKUP**

FDR InstantBackup is a separately priced member of the FDR family which enhances FDRABR.

For Archive backups, FDR InstantBackup enables support for the HSDM (High Speed Data Mover) feature available on some disk subsystems. HSDM allows FDR to backup and restore data in an internal compressed format. If your disk subsystem includes HSDM, you can invoke HSDM support by adding the DCT=YES operand on DUMP statements. [See Section 80.33](#) for details.

51.01 CONTINUED . . .

**ABR ARCHIVE
RESTORE**

Data sets which have been archived can easily be restored by a ABR batch job. Control statements specify the data set names which are required. ABR will search the Archive Control File backwards, in order to find the most recent copy of each requested data set. If a data set has been archived more than once, you can specify which copy you want. Multi-volume data sets are handled automatically, restoring all pieces of the data set.

The data from the ACF provides all the information necessary to locate and restore the requested files. You will usually use an option which allows ABR to dynamically allocate the backup files on tape or disk without JCL. ABR will sort the required backup files by volume serial and file number (for tape) so that each tape is mounted only once, even when many data sets are being restored.

**ABR AUTO-
RECALL**

When a data set is archived it can be marked as eligible for Auto-Recall ("Auto-Recall"). With Auto-Recall, the catalog entry for an archived data set is left in the MVS catalog, with a special indicator that indicates that it was archived. When an archived data set is referenced by batch JCL, dynamic allocation, or TSO, the ABR Catalog Locate Exit detects the indicator and automatically "recalls" the data set to disk before it is needed. Auto-recall makes the use of archived data sets transparent to the job or user, other than a short delay in execution while the data set is recalled.

TSO users normally have the option of bypassing the recall, and have a choice of waiting for the recall to complete (foreground) or doing the recall in a separate started task while they do other work (background). TSO users may also add recall requests (and some other ABR requests such as Archive) to a "remote queue" data set which is processed by a batch ABR job.

**TAPE MOUNT
MANAGE-
MENT**

ABR Archive can be used to implement the IBM concept of Tape Mount Management (TMM) where data sets which are normally on tape are directed to SMS-managed disk instead, and then moved to tape using ABR. This allows many such data sets to be stacked on tape and uses far fewer tape volumes. IBM has redbooks which describe TMM in detail (using IBM software, of course), but [Section 70](#) "SMS" contains additional information on implementing TMM with ABR Archive.

**ABR SUPER-
SCRATCH**

ABR Superscratch operates exactly like ABR Archive, but it does not backup the selected data sets, it simply deletes them. Superscratch can be used to delete data sets which you know will never be needed again (such as temporary data sets). Superscratch does not use the Archive Control File, since there are no backups to record and no way of recovering data sets which were deleted in error (unless you have another backup of the data, such as an ABR Volume Backup or Application Backup).

Simulation of Superscratch (see below) is recommended, to insure that valuable data sets are not inadvertently scratched.

There is a separate flag in the ABR Model DSCB to enable a volume for Superscratch.

Warning: Innovation strongly recommends that you enable the Superscratch flag in the ABR Model DSCB of a volume (using the FDRABRM utility) just before you execute the Superscratch job which processes that volume, and disabling the flag after the job is complete. This will prevent inadvertent Superscratch execution which could result in data loss. Also see the warning on Year2000 above.

SIMULATION

All functions of ABR Archive and Superscratch can be run in simulation mode, allowing you to verify that the correct data will be selected when run for real. Testing with simulation is especially important with Superscratch since an error in control statements may cause irreversible lost of data. Data sets Archived in error can always be restored.

51.01 CONTINUED . . .

ABR UTILITIES ABR includes a utility program, FDRARCH, for maintaining the Archive Control File. Since entries in the ACF are not automatically deleted when they reach their expiration dates, you must at least periodically run the REORG function of FDRARCH to purge those obsolete entries and make room for new ones. FDRARCH is documented in [Section 51.30](#).

FDRTSEL is a utility which can be used to move or copy ABR Archive backups. Among other things, you can use it to create a COPY 2 backup from a COPY 1, or to move a backup from disk to tape. FDRTSEL is described in [Section 60](#).

FDABRP ([Section 53](#)) is used for simple reporting on archived data sets, and FDREPORT ([Section 54](#)) can be used for more sophisticated reporting. The ABR ISPF dialog called SRS (Search, Report, and Services) can also display information on archived data; it is also in [Section 54](#).

51.01 CONTINUED . . .

**TAPE FORMAT
AND NAMING
CONVEN-
TIONS**

As indicated earlier, the backup files created by ABR Archive Backups are in standard DSF format. In this format, each backup file can contain data only from one DASD volume. So, if multiple DASD volumes are to be processed, ABR must create multiple backup files on the output tape or disk. Unless you restrict ABR to processing only a single disk volume in an ABR step, ABR will **always** create multiple backup files. If the amount of data archived from a given disk is excessive (over 4095 tracks by default), ABR may create multiple backup files, each with a subset of the data sets, in order to improve restore performance.

Since ABR must be able to uniquely name each tape file, and must be able to record the backup file in a way that it can easily be retrieved, ABR uses a special naming convention for the Archive Backup files. The name contains the disk volume serial, the date of the archive and a uniqueness character in case data is archived multiple times per day, so each ABR Archive Backup from a given disk volume will have a unique name.

The format is: **abrindex.Vvvvvvv.bnyydddx**
where:

abrindex	is the ABR hi-level index from the FDR Global Option Table. It is usually "FDRABR". It can be changed when ABR is installed but must NOT be changed once you start using ABR.
vvvvvv	is the volume serial of the disk volume from which the data sets were archived. ABR creates one or more backups file for each disk volume processed in a given ABR run.
n	is the copy number. ABR always creates COPY 1 and can optionally create COPY 2.
yyddd	Is the Julian date of the archive job (5 digits)
b and x	are qualifiers added to make the name unique if multiple Archive Backups are created for the same disk volume on the same Julian date. "b" will initially be set to "B" while "x" will cycle from A to Z and 0 to 9. When "x" has reached its maximum value (9), then "b" will reset to "D" through "J" while "x" cycles again. This allows for a maximum of 255 Archive Backups per disk per day. However, this name change occurs only if the name ABR is trying to create is already in the ABR catalog (see below); it may create multiple uncataloged backups with the exact same name.

Examples: FDRABR.VTSO002.B197147A (first COPY 1 created on 97.147 for volume TSO002)
FDRABR.VTSO002.D200350C (39th COPY 2 created on 2000.350 for volume TSO002)

Warning: you cannot override this naming convention, except for changing the high-level index during installation. As you will see in [Section 51.04 "Archive Job Control Requirements"](#) you must specify a data set name on the output (TAPEX) DD statements to satisfy MVS requirements, but that name will be ignored and the ABR-generated name used in its place.

The backup files created by ABR Archive Backup are usually not cataloged. For most such backups, all the information required is stored in the data set records in the Archive Control File. When restoring, information about the backup file to be read is obtained from the ACF so no entry in the ABR catalog is required. There are several exceptions:

- the ACF can only record 5 tape volumes for the backup file containing a specific data set. If the Archive backup of data from one disk volume required more than 5 tape volumes, the backup is cataloged to record the additional volume serials.
- In case tape management catalog control is being used for Archive tapes, the first file created on a tape by a given Archive job will be cataloged. Also, any time that the volume serial list of the current backup file is different from the serial list of the previous file (e.g., when the backup overflows the current tape volume and a new tape is mounted), that backup is cataloged.
- If EXPDT=99000 (catalog control for many tape management systems) is specified on the TAPEX or TAPExx DD statement, **all** backup files created on that DD will be cataloged.
- Archive backup files on disk are always cataloged.
- The ARCCAT=ALL operand can request that all archive backup files be cataloged.

51.01 CONTINUED . . .

**TAPE FORMAT
AND NAMING
CONVEN-
TIONS
(continued)**

If cataloging is required, the high-level index chosen (e.g., "FDRABR") must be assigned to an ICF catalog designated for ABR use, called the "ABR catalog". Since it is used mainly for Volume Backups, it is described in [Section 50.01](#).

As long as no more than 4 tape volumes are used for a given Archive Backup file, ABR will record a hardware "block id" in the Archive Control File for each backup file created on a cartridge drive. This supports IBM 3480, 3490, 3490E, 3590 Magstar, and all drives which emulate those IBM drives). During restore, this block id will be used to invoke a high-speed search at OPEN time to position directly to the beginning of each backup file. This can significantly reduce recall times, especially on high-capacity tapes such as the IBM 3590 Magstar and StorageTek Redwood and 9840.

**BACKUP
RETENTION
AND TAPE
MANAGE-
MENT**

Depending on the requirements for retention of Archived data sets established by your installation, you may choose to control the retention of ABR Archive Backups in several ways:

- You can let your tape management system expire backups based on a retention period or expiration date. This is known as "date control". If you are doing Archive Backups to disk, the retention/expiration is stored in the F1 DSCB of the backup data set, but the backup may not be deleted unless you run ABR Superscratch with the SELECT ABRBKUP statement against those volumes; FDRARCH can also delete expired disk backup data sets during Archive Control File reorganization.
- You can let ABR decide when to expire backups. When the REORG function of the utility FDRARCH is executed to reorganize an Archive Control File, it will delete obsolete entries, based on the expiration date recorded in the ACF or other criteria you specify. When the last data set entry which points to a given backup file is deleted, FDRARCH will uncatalog the backup file; if it is on disk it is also scratched. ABR has no formal interface to any tape management system, but if you use the "catalog control" feature of your tape management system the obsolete tapes are automatically scratched.

One way to implement the above is to tell FDRARCH to delete the entry for any data set which is no longer cataloged for Auto-recall in the system catalogs. If you Archive GDG generations, old generations are deleted when they "roll-off" from the GDG. Other data sets can be deleted from the Archive Backups simply by uncataloging the data set. This will also delete data sets which have been recalled, even if they were re-archived. When all the data sets on a given backup have been uncataloged and deleted, FDRARCH will uncatalog (and scratch on disk) the backup.

- If you have no tape management system, you must establish the necessary manual procedures for identifying and expiring ABR backups.

NOTE: if you have a tape management system from any software vendor, you should enable the TMS (Tape Management System) option in the FDR Global Option Table ([see Section 90](#)). The TMS option changes slightly the way that ABR handles files on tape to be compatible with restrictions of some such systems.

51.01 CONTINUED . . .

**BACKUP
RETENTION
AND TAPE
MANAGE-
MENT
(continued)**

When a Archive Backup is taken, an expiration date is assigned. It can come from several sources, in this order of priority:

- if the RETPD= operand is specified on the ABR DUMP statement, ABR uses that to calculate an expiration date which will be assigned to all COPY1 (TAPEX) backups created in that ABR step. If the RETPD2= operand is not specified but COPY2 (TAPEXX) backups are also created, the same expiration is assigned to those backups.
- If the RETPD2= operand is specified on the DUMP statement, ABR uses that retention period to calculate an expiration date which will be assigned to all COPY2 (TAPEXX) backups created in that ABR step.
- if the TAPEX or TAPEXX DD statement pointing to the backup tape contains the EXPDT= operand, the date specified is assigned to all backups written to that DD. For some tape management systems this value can also be a keyword; for example, EXPDT=99000 indicates "catalog control" to many tape management systems.
- If the DD statement contains the RETPD= operand, MVS uses that to calculate a expiration date, which ABR will assign to each backup written to that DD.
- If none of the above RETPD/EXPDT operands are given, the default retention period of 365 days (1 year) is used to calculate an expiration date which will be assigned to each backup.

By whatever means it is calculated, the expiration date is recorded in the Archive Control File record for each data set included in that Archive Backup file. If both COPY 1 and COPY 2 of a backup are being created, the expiration of each copy is recorded separately in case they are different.

The expiration date will be recorded in the tape labels of the backup file and will also be recorded by your tape management system, if you have one. If the expiration date is a real date (not a keyword such as 99000) then your tape management system will probably return the tape to scratch status on that date.

Your tape management system will usually scratch a backup tape (or multi-volume tape set) when every backup file on the tape has reached its expiration date. Since the backups created on a tape by a given ABR run will usually have the same expiration, this is not a problem. But there are options to place backups with varying expirations on the same tape, so you should beware of creating files on the same tape whose expirations are widely separated, since this may cause the tape to be retained longer than necessary. Since your tape management system and the Archive Control File each store the expiration of a backup file, you must insure that they do not contradict one another; if the expiration is changed in one, it must also be changed in the other.

If you create both COPY 1 and COPY 2 of your backups, you can use one copy for onsite recalls and the other copy for offsite recalls at a disaster/recovery site. Most tape management systems include vaulting support, allowing you to select tape to be sent offsite and returned when no longer required. Most users send COPY 2 offsite. Since the data set name used by ABR Archive Backup (described earlier) includes the copy number, you should be able to easily send COPY 2 backups offsite while retaining COPY 1 onsite.

51.01 CONTINUED . . .

**SMS-
MANAGED
VOLUMES**

On volumes that are managed by SMS (IBM's Storage Management Subsystem) you can choose whether ABR is to select data sets to Scratch or Archive by normal ABR criteria specified on ABR control statements, or by the attributes of the SMS management class associated with the data set. If the management class is chosen, and both SMS and non-SMS volumes are processed in the same ABR run, ABR criteria will be used for the non-SMS volumes and ignored for SMS volumes (with some exceptions). A complete explanation of Archive and Superscratch on SMS-managed volumes is in [Section 70](#).

There are advantages and disadvantages to each technique. ABR's Archive and Superscratch selection criteria are generally more flexible and offer more options than SMS, but may require many control statements if varying criteria are to be used for different data sets (usually controlled by data set name). SMS management classes can be associated with data sets at the time they are created as specified by the user or by installation ACS routines, independent of the data set name.

Even on SMS volumes, the decision over whether to create an archive backup on disk or tape (or both) is still controlled by ABR. However, if you are archiving based on SMS Management class attributes, you can optionally use the attributes of the SMS management class associated with each SMS data set to calculate expiration dates for COPY 1 on disk and COPY 2 on tape. Details are in "Archive Expiration" in [Section 70](#). This function allows you to use a global expiration for the tape file (such as EXPDT=99000 for catalog control), while assigning each data set a real expiration date based on its management class. With this option (the SMSEXPIRE= operand), tape management will record a global expiration for the whole backup file, using the rules above, but the ACF entry for each archived data set in that backup will record a specially calculated expiration. When all data sets reach their expirations, FDRARCH will uncatalog the backup file and make it eligible for scratch.

**ARCHIVE TO
DISK**

ABR Archive Backups can be directed to disk, instead of tape. This is usually used to provide a quick means of recalling an archived data set if it is needed in the first few days after it was archived; beyond that initial period, recall from tape is used. If you have an Automated Tape Library, such as those from StorageTek and IBM, you may want to do all of your archiving direct to tape, since the ATL can quickly mount required tapes and may be nearly as fast as recall from disk.

[Section 51.04](#) details the JCL requirements for directing Archive Backups to disk. The TAPEX DD statement will point to one or more disk volumes which are designated to hold Archive Backups. ABR will automatically allocate backup data sets on those disks, using the naming convention described earlier. These disk volumes are usually specified using the ABR POOLDISK feature, which allows ABR to monitor the space available on each volume and more efficiently use the disks. **ABR output disks cannot be SMS-managed.**

Backups of data sets to disk usually takes far less disk space than the original data sets occupied. The original interblock gaps are eliminated, only used tracks are backed up, and the backup can be compressed using FDR compression (described in [Section 51.03](#) and recommended for backup to disk).

You can use the FDRTSEL utility ([Section 60](#)) to move disk backups to tape. If a disk backup expires and is not moved to tape, the REORG function of the FDRARCH utility ([Section 51.30](#)) will delete and uncatalog obsolete disk backup files.

NOTE: when the Archive Backups are directed to disk, Innovation recommends that you specify the RTC=YES (or, if your disk subsystem supports the HSDM (High Speed Data Mover) option, DCT=YES). Either of these operands causes ABR to handle the output to disk efficiently. Failure to specify one of these operands causes ABR to use an I/O technique on the input data sets which is less efficient than that used when the outputs are on tape and causes. An Archive to disk to run longer than a similar Archive to tape.

51.01 CONTINUED . . .

**ARCHIVE
BACKUP
EXECUTION**

To execute ABR and perform Archive Backups, you must create ABR jobstreams and execute them at appropriate times. [Section 51.20](#) contains examples of such jobstreams which you can customize for your installation; all of the examples shown in this manual are also in the JCL library loaded as part of ABR's installation ([Section 90](#)). If you have an automation or scheduling product in your installation, you may want to use it to schedule daily ABR Archive Backups at appropriate times.

Here are several scenarios for ABR Archive Backups; you can pick the one that best meets your installation's requirements:

- **Archive direct to tape** - If the TAPEX DD points to tape, then the COPY 1 backup will be on tape. If you include the optional TAPEX DD statement, COPY 2 will be created on tape as well. You will probably want to assign a long-term retention to both copies. You might want to send COPY 2 to offsite storage so that archived data sets can be recalled at a disaster/recovery site.
- **Archive to disk only** - if the TAPEX DD points to disk volumes (as described above), ABR will allocate the COPY 1 backup on disk. Since the disk space allocated to ABR is limited, you will probably want to specify a fairly short retention period (e.g., 7 days, 14 days, 31 days) for the disk backups. During that period, data sets can be recalled quickly since no tape mount is required. You can execute FDRTSEL ([Section 60](#)) to move the COPY 1 backups to tape with a longer retention as they approach their expiration date on disk; FDRTSEL can also be used to create a COPY 2 on tape from the COPY 1 on disk.
- **Archive to disk and tape** - if the TAPEX DD points to disk, and the TAPEX DD points to tape, ABR will allocate the COPY 1 on disk and the COPY 2 on tape. You will probably want to specify a fairly short retention period for COPY 1 and longer retention for COPY 2 (during recall, ABR will automatically switch to COPY 2 if COPY 1 has reached its expiration date as recorded in the Archive Control File). You may want to execute FDRTSEL to move the COPY 1 backups to tape with a longer expiration so that you are not left with a single backup of the archived data.

If you have an Automated Tape Library (ATL), you may want to put COPY 1 to ATL tapes, since they will be automatically mounted when required, providing recalls that are almost as fast as recall from disk.

For data set recovery from Archive Backups, there are several alternatives:

- you can permit users to restore data sets directly by submitting their own RESTORE TYPE=ARC job. This may require mounting backup tapes for each data set required. If FDR security is enabled, users will need UPDATE or ALTER authority to the data sets.
- you can allow users to add requests to a remote queue (see details in [Section 51.03](#)). In this case Operations (or your automation software) must submit a RESTORE TYPE=ARC job at regular intervals (e.g., every hour, every 4 hours) to process any requests on the queue. If multiple data sets are being restored, this minimizes tape mounts. If FDR security is enabled, users will need UPDATE or ALTER authority to the data sets.
- You can require that users communicate their data set restore requests directly to someone in the Data Center, who will submit the required RESTORE TYPE=ARC job.
- If the Auto-Recall option was in effect when the data set was archived, users can recall archived data sets simply by referencing them from TSO or a batch job. The data set will be automatically recalled from the most recent Archive Backup tape on which it appears. Data sets which can be Auto-Recalled must all be recorded in a common Archive Control File (ACF), so one ACF is usually used for all Archive Backups.

51.01 CONTINUED . . .

**ARCHIVE
BACKUP
MANAGE-
MENT**

Unlike Volume Backups, which are usually kept for a fairly short period of time (so many weeks, or so many generations), Archive Backups may be kept for a long time. The default retention is 1 year, but many installations make it longer. If you do archiving frequently, you could end up with hundreds or thousands of tape volumes containing active Archive Backups.

The problem is compounded because much of the data on those backups may be obsolete. For data sets which have been deleted from the Archive Control File by FDRARCH (for expiration date, uncataloged, etc.), pointers to the backups of those data sets no longer exist. But as long as one archived data set is still recorded pointing to an Archive Backup file, that file will be retained. Archive creates multiple files on tape, and may create multi-volume aggregates, and most tape management systems will not scratch any tapes in the aggregate until every file is eligible for scratch.

So, as Archive Backup tapes get older, the chances are that the amount of active recallable data on the tapes will decrease.

This problem is addressed by the ABR utility FDRTSEL, described in [Section 60](#). FDRTSEL can copy your old Archive Backup files to new tape. Only the active files on an input tape are copied; for example, if a tape contains 10 files, but the Archive Control File has active records for only 3 of those files, only those 3 are copied to the new tape.

FDRTSEL can take files from various input tapes and combine them on one output tape, reducing the amount of tape required.

Normally, FDRTSEL will copy the entire contents of a tape file it selects, possibly copying data for data sets which are no longer active in the Archive Control File. If the ARCEDIT option of FDRTSEL is used, it will copy only the data belonging to data sets still recorded in the ACF; this reduces the amount of tape required for output even further.

Regular use of FDRTSEL, perhaps monthly, will help keep the number of tape devoted to Archive Backups to a minimum. However, its effectiveness is dependant on your techniques for purging data sets from the ACF. For example, if you don't purge uncataloged data sets or expire based on SMS management class attributes, then all the data in a given backup file will always expire on the same day and ARCEDIT will be useless.

51.02 DATA SET SELECTION

**ARCHIVE
BACKUP AND
SUPER-
SCRATCH**

For an Archive Backup (DUMP TYPE=ARC) or SUPERSCRATCH (DUMP TYPE=SCR) operation, ABR uses the following rules for data set selection:

- If SMSMANAGE=YES is specified on the DUMP statement, SMS-managed data sets will be selected for ARCHIVE or SUPERSCRATCH based on the attributes of their associated SMS management class as described in [Section 70](#). If a data set is selected by its SMS attributes, there is no option to exclude it; it will be archived or scratched. If it is not selected by SMS, you can optionally let ABR SELECT statements determine if it should be selected (SMSCOMMAND=YES).
- If the ABR ARCHIVE PROTECT LIST (for TYPE=ARC) or SCRATCH PROTECT LIST (for TYPE=SCR) is enabled ([see Section 90](#)), all data sets which match entries in the protect list will be excluded.
- If the ABR job includes SELECT statements, those statements are compared to every data set on every volume processed by ABR. The SELECT statement contains one set of operands, called Data set Selection Criteria, used to determine if the SELECT matches a given data set, such as DSN=, VOL=, DSORG=, etc. If the data set matches **all** of the Data set Selection Criteria specified, it will be selected by that SELECT statement but not necessarily archived. Otherwise, the data set will be compared against the following SELECT statements.
- SELECT statements can contain a second set of operands, known as Archive Selection Criteria. They include attributes such as IFNOTCAT (archive if not cataloged) and ADAYS=nn (archive if not opened in the last nn days). If the SELECT which matched a particular data set contains one or more of the Archive Selection Criteria, it will be archived if it matches **any** of those criteria, and it will remain on disk if not. If the SELECT contains none of the Archive Selection Criteria operands, the data set will be unconditionally archived.
- The DUMP statement can optionally contain additional General Archive Selection Criteria operands, similar to those on the SELECT statement. If a given data set was not selected by any SELECT statement, but the DUMP statement contained criteria operands, the data set will be archived if it matches **any** of those criteria, and it will remain on disk if not. This is also true if there are no SELECT statements.

ABR will read the VTOC of every disk volume in its volume list ([See Section 51.03](#)) and execute the selection process described above against every data set it encounters.

If the ADATE= or ADAYS= operand is specified, ABR will examine the last reference date that IBM stores in the Format 1 DSCB of every data set when it is opened. If it is less than ADATE or today's date minus ADAYS, the data set will be archived or scratched.

ABR allows users to request that certain data sets be included in the next archive, if the volumes containing those data sets are processed, using a "remote queue". Details are in [Section 51.03](#).

51.02 CONTINUED . . .

**ARCHIVE
BACKUP AND
SUPER-
SCRATCH
(continued)**

This selection process can be confusing, so here are some examples of the different ways that you could select data sets for archive:

- **General Archive Selection Criteria on the DUMP statement and no SELECT statements.**

e.g.,

```
DUMP TYPE=ARC,ADAYS=15,IFNOTCAT
```

all data sets on selected volumes will be compared to the criteria. Those that are uncataloged or unused in the last 15 days are archived or scratched.

- **No criteria on the DUMP statement and SELECT statements with Archive Selection Criteria, e.g.,**

```
DUMP TYPE=ARC
SELECT ALLDSN,ADAYS=15,IFNOTCAT
```

this is equivalent to the previous example.

```
DUMP TYPE=ARC
SELECT DSN=XYZ**,ADAYS=15,IFNOTCAT
```

all data sets on selected volumes will be compared to the SELECT. Those starting with XYZ are controlled by this SELECT. Those that are uncataloged or unused in the last 15 days are archived or scratched.

- **No criteria on the DUMP statement and SELECT statements with no criteria, e.g.,**

```
DUMP TYPE=ARC
SELECT DSN=XYZ**
```

all data sets on the selected volumes that start with XYZ will be unconditionally archived.

- **General Archive Selection Criteria on the DUMP statement and SELECT statements, e.g.,**

```
DUMP TYPE=ARC,ADAYS=30,ONLYVOL
SELECT DSN=ABC**,VOLG=TSO
SELECT DSN=XYZ**,ADAYS=15,VOLG=DB
```

all data sets starting with ABC on TSO volumes are unconditionally archived. All data sets starting with XYZ on DB volumes are archived if they have not been used in 15 days. All other data sets on the selected volumes are archived if they have not been used in 30 days. This combination allows you to specify general criteria for most data sets while defining specific criteria for special groups of data sets.

ICF VSAM

ABR can only Archive or Superscratch ICF VSAM files by cluster name. You cannot select individual components. Once ABR selects a cluster from a volume, all components of the cluster on that volume (including alternate index components) are dumped.

ICF catalogs cannot be archived or scratched.

Multi-volume VSAM (clusters with components or parts of components on multiple disk volumes) will be scratched only if all the volumes on which the cluster exists are processed in the same ABR step, and the cluster is selected from all of its volumes.

IBM stores the last reference date only in the Format 1 DSCB on the first volume of the data component of the base cluster ([See Section 80.13](#) for more details). If it is selected, all components on the volume containing that DSCB are selected. For multi-volume VSAM, if the cluster is selected on its first volume by ADAYS= or ADATE=, this will not cause it to be backed up from the secondary volumes, nor to be scratched. The member \$\$VSAM in the FDR ICL (Installation Control Library) discusses this topic in more detail, and provides a sample jobstream to ARCHIVE multi-volume VSAM by last reference date.

51.02 CONTINUED . . .

**EXCEPTIONS
TO
AUTOMATIC
ARCHIVING**

The following data sets will **never** be selected for Archive or Superscratch:

- For Archive, data sets in the ARCHIVE PROTECT LIST, if the protect list is enabled in the FDR Global Option Table (ARCPROT/PROTECT). For Superscratch, data sets in the SCRATCH PROTECT LIST if enabled (SCRPROT/SCREXCL).
- Data sets marked 'do not archive', i.e. data sets for which OPTIONS=AD or ND or EX have been set with program FDRABRM.
- Data sets on volumes whose ABR Model DSCB indicates 'disabled for Archive' for TYPE=ARC or 'disabled for Superscratch' for TYPE=SCR. These flags are set by program FDRABRM or the ABR TSO/ISPF panels.
- Data sets which begin with the ABR prefix (normally 'FDRABR' unless overridden in the FDR Global Option Table). However, if SELECT ABRBKUP is specified then ABR will select ABR backup data sets which are expired.

The following data sets will not be selected by General Archive Selection Criteria on the DUMP statement or by a SELECT ALLDSN statement, but can be selected by other SELECT statements with more specific DSN= operands, including DSN=**:

- Data sets prefixed by SYS1.
- OS PASSWORD data set.
- SYSCTLG data set (OS CVOL catalog).
- Temporary data sets. If you want to select temporary data sets, specify SELECT TEMP preceding any SELECT ALLDSN or any other SELECT that would select the same data sets. DSNENQ=USE and CRDAYS=1 are recommended to be sure that temporary data sets in use will not be selected.
- Data sets found to be in-use when using the data set enqueue option (DSNENQ=), unless ENQERR=PROCESS is specified.
- Data sets with a last reference date of zero (00000) if data set selected by ADAYS= or ADATE=.
- Non-ICF VSAM clusters.
- Model DSCBs (data sets with no extents).

If a data set is selected by a SELECT statement (and not excluded by a preceding EXCLUDE statement), it will be archived or scratched unless it matches one of the first four bullets in the above list. If a data set is selected by a SELECT ALLDSN statement, it will be archived or scratched unless it falls into any of the categories listed above.

**SELDOM-
OPENED DATA
SETS**

There is one situation which "fools" Archive Backup into archiving data sets which are really in use. The Last Referenced Date in the Format 1 DSCB of a data set is updated **only when the data set is actually opened**. If a job has a DD statement referencing the data set but does not open it, the date is not updated. This may cause the data set to be archived, since it appears that it has not been recently used, yet it will be immediately recalled the next time that the job is run. This is useless effort, and will delay the start of the job while the recalls are done.

The most common occurrence of this involves CICS, specifically a test CICS region which may have DD statements pointing to test data files for a variety of applications. CICS may not open those test files unless the associated application is being updated or tested, so it may not open some of those files for months. If ABR archives them, they will be recalled the next time that the test CICS region is started.

If you can identify such data sets, you may want to add them to the ARCHIVE PROTECT LIST so that they will never be archived.

51.02 CONTINUED . . .

ARCHIVE ABR will normally backup the Archive Control File (ACF) at the end of each Archive Backup step.
CONTROL FILE Since the ACF is necessary for the restore of archived data sets, it is backed up immediately for
BACKUP safety. For tape output, this backup becomes the last file on the last active output tape. For disk, it will be allocated on one of the output disks.

Two types of backup are available for the Archive Control File:

- The default (ARCBACKUP=ABR) causes an ABR incremental Volume Backup of just the Archive Control File; essentially it does an ABR DUMP TYPE=DSF selecting only this one data set. It does become part of the Volume Backups for the volume the ACF resides on, so you must be doing Volume Backups of that volume. It will increment the Cycle number for that volume and use the naming convention of ABR Volume Backups as described in [Section 50](#). It may result in a full-volume backup. If it is necessary to use ABR to recover that volume from backups (RESTORE TYPE=FDR), ABR will have to read all of the regular incremental backups plus all the incrementals containing only the ACF, even though it will use only the most recent. For these reasons, ARCBACKUP=ABR is not recommended (remember that it is the default).
- The second type of backup (ARCBACKUP=DSF) causes ABR to use FDRDSF directly (not through the ABR incremental system) to backup the ACF. The backup file will have the same data set name as the ACF itself, except that the index level of 'ARCHIVE' will be changed to 'ARCBKUP' for TAPEX and 'ARCBKU2' for TAPEXX (if duplicate backup). This cataloged name will only keep track of the most current backup.

For example, if the ACF is named "FDRABR.ARCHIVE.RECALL" the backup files containing the DSF backup of that data set will be named "FDRABR.ARCBKUP.RECALL" and "FDRABR.ARCBKU2.RECALL".

If the Archive job is backing up to disk, ABR will default to not backing up the ACF (ARCBACKUP=NO). This default should not be overridden.

51.02 CONTINUED . . .

**RESTORE
FROM
ARCHIVE
BACKUP**

In most cases, archived data sets are restored automatically by ABR Auto-Recall, described in detail later in [Section 51](#). But you can manually restore archived data sets at any time.

To restore a data set from Archive Backup, you must execute an ABR "RESTORE TYPE=ARC" job and provide one or more SELECT statements to identify the data sets required. The job can also specify the Archive Control File (ACF) to be used for the restore, but in most cases there will only be one ACF for all archived data sets in your installation.

Since the ACF is organized in chronological order based on Archive date, ABR will search for the data sets requested by reading it in reverse order. If a data set has been archived more than once, this will find the most recently archived version. If you want a different version, you can specify ADATE= (date of archive) or OLDBACKUP= (relative archived version) to tell ABR which version to restore.

If a multi-volume data set was archived, each volume will have a separate entry in the ACF, but each will be flagged to indicated it was part of a multi-volume entry. ABR will continue to search the ACF until it finds all of the entries belonging to the data set.

If NEWNAME/NEWGROUP/NEWINDEX is specified, the data set will be restored using the new name; otherwise it will be restored using its original name. If NVOL= is specified, the data set will be restored to that specified volume serial (or the first volser specified); otherwise the catalog will be searched for the output data set name (either the original name or the new name). If found, the data set will be restored to the volume serial specified in the catalog; if not found, the restore will be attempted to the volume serial from which the data set was archived.

The restore process is identical to that of FDRDSF. Please review [Section 20.01](#) for details on allocating, restoring, and cataloging data sets during the restore. [Section 51.08](#) contains details on output volume selection for ABR restores.

ABR will default to restoring from COPY 1 of each Archive Backup required unless COPY 1 was not created or has expired. COPY2 will then be selected if it exists. If TAPE=EXP is specified, ABR will ignore the expiration date test and attempt to restore from COPY 1 even if it is expired.

The **ABR Restore Protect List**, documented in [Section 90](#), can list data sets or groups of data sets which are not to be restored, similar to EXCLUDE statements. Any attempt to restore those data sets from Volume Backups or Archive Backups will fail (except for full-volume recovery from Volume Backups).

51.03 PROCESSING OPTIONS AND REQUIREMENTS

ABR OPERATIONS The first control statement in the ABR input must be a DUMP or RESTORE statement. It may be followed by SELECT and optionally EXCLUDE statements to specify data sets which must be backed up. It may also be followed by MOUNT statements to specify the DASD volumes or SMS storage groups to be processed in this execution of ABR (DISKxxxx DD statements can also be used to specify volumes as described in [Section 51.04](#)). ABR will accept up to 100 SELECT/EXCLUDE/MOUNT statements in a single execution, unless that limit is overridden.

ABR ARCHIVE AND SUPERSCRATCH OPTIONS The DUMP statement contains a TYPE= operand to specify the type of backup. In each case, volumes to be processed in the ABR step are identified by MOUNT statements and/or DISKxxxx DD statements. These 2 values will invoke ABR Archive Backup or Superscratch:

A. TYPE=ARC

ABR Archive Backup will be executed for every selected volume. Selected data sets will be backed up, recorded in the Archive Control File, and scratched from disk.

B. TYPE=SCR

ABR Superscratch will be executed for every selected volume. Selected data sets will be scratched. No backup will be taken. No recovery is possible unless you have ABR Volume Backups or other backups of the data sets.

In both these options, ABR builds a list of the DASD volumes to be processed. ABR's list of disk volumes to be processed comes from several sources: if DISKxxxx DD statements are present in the ABR step ([see Section 50.04](#)) the named volumes are added in the order they appear. If MOUNT statements are present in the input, the named volumes are added in the order of their UCB addresses. If ONLVOL is specified on the DUMP statement, any volumes named on SELECT statements are included. If ONLINE is specified on the DUMP statement, all online volumes are included.

For Superscratch, the volumes are processed in their table order. A TAPEX DD DUMMY is required but is not used.

If the Archive step JCL contains a single TAPEX DD statement, ABR will select the first disk volume on its list and do the requested backup, creating file 1 on that output tape. When complete, it selects the next volume and creates file 2 on the tape, etc., until all volumes are processed. If the amount of data selected from one volume exceeds a limit (4095 tracks by default), ABR will create multiple backup files for that volume, each containing a subset of the data sets, to improve restore performance.

However, you may have up to 9 TAPEX DDs in the Archive step. If you have more than one, ABR will select the first "n" volumes on its list and assign them to the "n" tape drives, executing the backups in parallel with internal subtasking, and creating file 1 on each output tape. As each backup completes, ABR selects the next disk volume and assigns it to the tape which just became idle, creating the next file number on that tape. Since the elapsed time of an Archive run is usually not important (these are inactive data sets, after all), a single TAPEX DD is usually sufficient. See "Memory Requirements" below for limitations on the use of multiple TAPEX DDs.

51.03 CONTINUED . . .

**MEMORY
REQUIRE-
MENTS** See "Memory Requirements" in [Section 50.03](#)

**COMPRESS
OPTION** The dump can be instructed to compress the data on the sequential backup file using Innovation's own proprietary software compression algorithm.

It is **not recommended** for backups to tape attached by ESCON or FICON channels because of the speed of the channel. All ESCON/FICON-attached tape drives include hardware compression which will provide savings similar to that of FDR compression. Software compression will be ignored for backups created with the HSDM disk hardware option (DCT=YES) since the data is pre-compressed.

NOTE: all FDR restores will automatically recognize a compressed backup file and decompress it. No special option is required to restore a compressed backup.

**DUPLICATE
TAPE OPTION** ABR has an option to create a duplicate or second copy of the backup tape during dump processing. The primary copy is called COPY 1 and the duplicate is called COPY 2. The copy number becomes part of the name of the backup data set as described in [Section 51.01](#).

While dumping a disk to a TAPEX DD statement, the duplicate backup will be written to the TAPExx DD statement (same "x" value twice) if it is present. You may have TAPExx DDs for some TAPEx DDs and not for others in the same step but this is not recommended with ABR since you cannot predict which disks will be written to which tapes.

Memory requirements do not increase with the use of the duplicate tape option.

51.03 CONTINUED . . .

SECURITY Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system which supports the SAF (Security Authorization Facility) interface, such as ACF2 or TOP SECRET, you can enable the ALLCALL option. For ABR Archive Backup and Superscratch this results in these security checks:

- for Archive Backup (DUMP TYPE=ARC) and Superscratch (DUMP TYPE=SCR), ABR will check to see if your userid has at least ALTER authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority or the input volume is not protected by your security system, then ABR will check if you have at least ALTER authority under the DATASET security class to every data set selected from the input volume. Any data sets to which you are not authorized will be bypassed with an error message.
- for manual data set restores (RESTORE TYPE=ARC), ABR will check if you have at least UPDATE authority under the DATASET security class to every data set restored. Any data sets to which you are not authorized will be bypassed with an error message. If an output data set must be allocated, the operating system will check if you have CREATE/ALLOCATE authority for the data set (this is done even if ALLCALL is not enabled). The user may need READ authority to the Archive Backup files on tape or disk (even if ALLCALL is not enabled).
- for Auto-Recall of archive data sets, the above restore security checks may apply. However, it is possible to configure ABR so that only READ authority is required; this allows a job or TSO user to recall a data set which it will use as input.

Innovation recommends that volume security rules be established for the userids which will be executing ABR Archive Backups and Superscratch. This will minimize the security checking overhead as well as providing clearcut control over who can execute the backups. Under RACF, volume security is controlled by the DASDVOL security class, and other security systems have similar facilities (see their documentation). Under RACF, it is also possible to execute under a userid with OPERATIONS authority, which automatically grants DASDVOL and DATASET authority.

51.03 CONTINUED . . .

DATA SET ENQUEUE OPTION You can request, via the DSNENQ= operand, that each data set being archived, scratched or restored be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name or SYSDSN and a minor name of the data set name itself, for each selected data set found in the VTOC of the input disk; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. Note that FDR cannot tell if the data set is being used for input or output. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations.

Optionally you can request that inactive data sets be enqueued to ABR during the backup or superscratch, to insure that no other job or TSO user can access the data set until ABR is done with the volume.

For Archive Backups and Superscratch, in-use data sets will be bypassed. You will receive the FDR158 message documenting the active data sets as a warning, but it will not be considered an error; the step will end normally if no other errors occur.

For restores, ABR will attempt to enqueue any data sets that it allocates on the output disks, to insure that no other task tries to use them until the restore is complete, but if the ENQ fails, the data set is still restored. But for existing data sets, if the ENQ fails, the restore will be bypassed.

The DSNENQ= operand has 3 possible values:

USE – data sets will be enqueued for the duration of the backup or restore on this disk volume. For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option. **This is the default for TYPE=ARC restores.**

TEST – data sets will only be tested to see if they are enqueued to another task at the time that the backup or restore on this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the backup or restore is proceeding.

NONE – No data set ENQ will be issued. **This is the default for Archive Backup and Superscratch.**

NOTE: If a data set name appears in a DD statement with DISP=SHR within the ABR job (not necessarily in the ABR step), and you specify DSNENQ=USE or TEST, ABR will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the ABR job ends.

CAUTION: *This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.*

Recommendation: DSNENQ=USE is strongly recommended to prevent ABR from attempting to scratch data sets which are actually in use by another task. It also prevents Archive from backing up data sets which it can't scratch. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

51.03 CONTINUED . . .

**VTOC
ENQUEUE
OPTION**

ABR also supports, via the ENQ= operand, an ENQ on the VTOC of every volume being dumped. For shared DASD, it can also invoke a hardware RESERVE on the volume during the FDR operation.

The VTOC is protected by an ENQ with major name SYSVTOC and a minor name of the volume serial. This ENQ is held by any task doing updates to the VTOC, including allocation of new data sets, extension of data sets to new extents, and scratching of existing data sets. This ENQ is normally of short duration, just for the few seconds necessary to update the VTOC, so if the ENQ is currently held by another task, ABR will wait for it to be released.

The SYSVTOC ENQ does **not** prevent access to existing data sets on the volume; it only insures that no other task is updating the VTOC while ABR is processing it. VTOC changes during a backup could result in an invalid backup.

For disks shared with another MVS system or LPAR, ENQ=RESERVE requests that, in addition to the ENQ described above, a hardware RESERVE is done on the volume. RESERVE will prevent any system from doing any I/O on the volume, except for the system that ABR is running on where only the ENQ protection applies. If you have a cross-CPU ENQ facility, such as GRS or MIM, you may be able to convert the RESERVE into a cross-CPU SYSVTOC ENQ and allow access to the volume during the operation (look up SYSVTOC in the documentation for your product).

Use ENQ= to prevent other tasks from making changes to the VTOC during the operation. ABR Archive Backups and Superscratch default to ENQ=ON (ENQ on same CPU). If you are executing ABR on shared DASD, specify ENQ=RESERVE. Member ENQ in the FDR ICL (Installation Control Library) has more information on VTOC ENQs

**STEP
TERMINATION**

If no errors occur during the execution of ABR, the ABR jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

- Some errors are critical. The jobstep ends immediately with a user ABEND.
- Some errors are critical only to a particular operation. For example, during a backup, some errors cause the backup of a particular disk to terminate immediately, but ABR may continue and attempt to backup other disks requested in the same step.
- Some errors are non-critical and the messages are warnings only. ABR will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with return code 12 to call your attention to the errors. Remember that RC=12 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer to get a different return code or a U0888 abend on a non-terminating error, the ABRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice or ABEND ([see Section 90](#)).

51.03 CONTINUED . . .

REMOTE QUEUE ABR includes a facility, called "remote queue", where users can request that certain data sets be included in the next Archive Backup, or that certain data sets be restored from Archive Backups. Requests can be added to the queue by TSO users using the ABR ISPF dialog; there is also a batch utility (FDRABRUT) for adding requests.

For backups, depending on installation options, a backup request to the remote queue can simply turn on a flag in the F1 DSCB of the data set, or it may add a statement to a remote queue data set requesting the backup. For restores, a remote queue data set is always updated. You will find more information on the setup of the remote queue data sets in [Section 90](#) (Installation). They are honored only if you setup the data sets properly and include the proper DD statements in ABR jobs to process them.

If a remote queue data set is used for backups, a special DD statement must be included in the backup JCL ([See Section 51.04](#)) pointing to the data set. ABR will add the requests (which act like SELECT statements) to the ABR control statement input.

If a remote queue is used for restores, you will probably want to run a special ABR data set restore job at intervals (perhaps several times a day) just to process the remote requests so that the users do not have to wait an excessive time for their data sets to be restored. A special DD statement must be included in the restore JCL ([See Section 51.07](#)) pointing to the data set. ABR will add the requests (which act like SELECT statements) to the ABR control statement input, which normally consists only of a RESTORE TYPE=ABR statement.

DYNAMIC ALLOCATION ABR will dynamically allocate disk volumes as needed. There is no need to provide DISKxxxx DD statements for volumes to be processed by ABR. As long as the required volumes are online, ABR can dump from and restore to any required disk volumes.

For restore operations, ABR will dynamically allocate each required backup tape if the DYNTAPE option is specified. If required backups are on mixed device types (such as 3480s and 3490Es), ABR will automatically mount each tape on the proper device type. For automated tape libraries, a drive in the proper library will be allocated. If you have multiple tape libraries, you may need to enable the DYNDEALC option in the FDR Global Option Table (ISPF panel A.I.4.4, [see Section 90](#)).

For data set restores, ABR will sort the list of backup files required and mount the backup tapes in an order which minimizes the amount of tape movement required.

VOLUME THRESHOLDS ABR can optionally process volumes for Archive and Superscratch based on allocation thresholds (percentage of the tracks on the volume allocated to data sets). For example, you may choose to process certain volumes only if they are more than 90% allocated, and bypass other volumes that are below this threshold. The purpose of this is to reduce ABR overhead and run time by processing only those volumes which need to have space freed up for the allocation of new data sets.

For non-SMS volumes, a high and low threshold (as a percent from 0 to 100) can be stored in the ABR model DSCB. The FDRABRM utility statements MAINT, REMODEL, and ABRINIT accept operands to set these thresholds ([See Section 50.40](#)), and the ISPF A.I.8 panel for maintaining ABR volume options can also display and set them ([See Section 50.46](#)). For SMS-managed volumes, the SMS storage group associated with each volume contains the thresholds (although the thresholds in the ABR models on those volumes can optionally be used instead).

With thresholding, ABR will still build a list of volumes by its usual techniques (DISK DD statements, MOUNT statements, and/or the ONLINE/ONLVOL operands), but it can be instructed to bypass any volumes under their high threshold, or under their low threshold, or under a threshold specified in the ABR run. By default, no threshold tests will be done and ABR will process all volumes in its list.

Even with thresholding, once a volume is selected for processing, ABR will scratch or archive all eligible data sets from the volume. This is different from the implementation of thresholds in DFHSM. DFHSM will stop selecting data sets when the allocation percentage on the volume has been reduced below its low threshold, where ABR will not. ABR can afford to do this because our backup technique is very efficient.

51.04 ARCHIVE AND SUPERSCRATCH JOB CONTROL REQUIREMENTS

The following Job Control Statements are required to perform Archive Backup and Superscratch functions.

**STEPLIB or
JOBLIB DD
STATEMENT**

If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.

**EXEC
STATEMENT**

Specifies the program name (PGM=FDRABR), region requirement (REGION=, [see Section 51.03](#)), and optional PARM= operand.

If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid DUMP statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR    EXEC    PGM=FDRABR , PARM=' DUMP  TYPE=ARC , ADAYS=15 '
//FDR    EXEC    PGM=FDRABR , PARM=' DUMP  TYPE=SCR /  SELECT  DSN=A . B . C '
```

If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

**SYSPRINT DD
STATEMENT**

Specifies the primary output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**SYSPRINx DD
STATEMENT**

Specifies the secondary output data set for messages related to the matching TAPEx DD statement. A SYSPRINx is required for each TAPEx present in the step. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**ABRMAP DD
STATEMENT**

Specifies the output report data set if SIM or PRINT=ABR is specified. It is optional; if not specified the SIM report will be printed on the SYSPRINx DD statement. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

**FDRSUMM DD
STATEMENT**

(Optional) if present, ABR will write one-line messages for each volume dumped, giving result codes, elapsed time, and byte counts. Usually a SYSOUT data set. FDRSUMM is used only if RTC=YES or DCT=YES is specified on the DUMP statement. Note that a volume may appear more than once in FDRSUMM if the MAXBTRKS= value causes the volume to be processed more than one.

**SYSUDUMP DD
STATEMENT**

Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

**ARCHIVE DD
STATEMENT**

Specifies the ABR Archive Control File. This DD is required for DUMP TYPE=ARC unless the DYNARC operand is specified. The standard name for this data set is FDRABR.ARCHIVE, but you can use any name as long as it has an index level of "ARCHIVE" somewhere in the name. ABR uses this data set as an inventory and catalog file for all data sets that have been archived from disk volumes. Specify DISP=SHR since ABR internally serializes access. On SIMULATION or TYPE=SCR operations this DD is not required.

```
Example : //ARCHIVE DD DISP=SHR , DSN=FDRABR . ARCHIVE . MAIN
```

51.04 CONTINUED . . .

**DISKxxxx DD
STATEMENTS**

Optionally specifies disk volumes to be processed by ABR. The format will be:

```
//DISKxxxx DD UNIT=unitname,VOL=SER=volser,DISP=OLD
```

"xxxx" may be 1-4 alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) characters, "unitname" is either a generic name, such as 3390, or an esoteric name assigned during your I/O configuration, such as DISK or SYSALLDA, and "volser" is the volume serial of the disk volume (if an esoteric unit name is used, the volume serial must be mounted on a disk unit which is part of that esoteric). You may use either DISP=OLD or DISP=SHR; it makes no difference.

The DISKxxxx DD may also specify multiple volume serials, up to 255, in the format:

```
//DISKxxxx DD UNIT=unitname,DISP=OLD,  
// VOL=SER=(volser1,volser2,...)
```

All the volumes must be of the same type (e.g., 3390) and have the same capacity (e.g., 3390-3).

ABR will build a list of the volume serials from all of the DISKxxxx DDs found in the ABR step JCL; these volumes will be processed in the order that the DISKxxxx DDs are found. A maximum of 256 disk volumes can be processed in one execution of ABR; if more disk volumes are required, MAXDD= must be specified on the DUMP statement.

However, DISKxxxx DD statements are not required and there are easier alternatives. If the ONLINE operand is specified on the DUMP statement, all online disk volumes will be processed automatically. If MOUNT statements are included in the ABR input, ABR will automatically process the indicated volumes (by volser, volume serial group, or SMS storage class). If the ONLVOL operand is specified, any volumes specified by SELECT statements (from the VOL= or VOLG= operands, or selected from the catalog by the CATDSN= operand) will be automatically processed. These volumes will be added to the list of volumes specified by DISKxxxx DDs (if any). Both ONLINE, ONLVOL and MOUNT will select volumes in the order of their device addresses (UCB address order), so you may still want to use DISKxxxx DDs to specify the order of processing for certain volumes.

NOTE: If ABR encounters multiple DISKxxxx DD statements pointing to the same disk volume serial, ABR will process the volume once using the first DDNAME encountered. DISKONLx is a reserved DDNAME, used by ABR for allocation of volume that do not have DISKxxxx DDs.

51.04 CONTINUED . . .

**TAPEX DD
STATEMENT**

Specifies the output device on which backup files are to be created. "x" may be any single alphabetic (A-Z), numeric (0-9) or national (@ # \$ in the US) character. A maximum of 9 such TAPEX DDs may be present in the ABR step JCL. ABR will start an internal dump subtask for each TAPEX DD present, each subtask processing one disk volume from ABR's list of disk volumes. If only one TAPEX is provided, ABR will process disk volumes one at a time; if multiple TAPEX DDs are present, ABR will process up to 9 disks in parallel (however, region requirements may limit the number of concurrent subtasks; [see Section 51.02](#)). TAPEX DDs that specify DUMMY or DSN=NULLFILE will be ignored except for simulation and Superscratch.

Tape Output:

If outputting to tape or cartridge, the TAPEX DD must specify:

DSN=	a data set name is required by MVS, but it will be overridden by ABR at OPEN time, so any non-temporary name is acceptable. The name you specify will not be used by ABR, but MVS will do an exclusive ENQ on this name at job initiation so each ABR job should use unique names.
UNIT=	specify a generic (e.g., 3490) or esoteric (e.g., CART) name to allocate the type of tape drive desired. If you have sufficient tape drives available, specifying a unit count of 2 (e.g., UNIT=(3480,2)) may reduce elapsed time (especially on 3480 cartridge drives).
VOL=	specify a volume count, e.g., VOL=(, , 255), to prevent ABR from abending if more than 5 tape volumes are required. If no volume serials are specified, ABR will call for scratch tapes; this is recommended; however, you can specify up to 255 tape volume serials if you need to select the output volumes.
LABEL=	you may want to specify RETPD= or EXPTD= to identify the expiration date of the backups (retention periods may also be specified on the DUMP statement). ABR stores this expiration date in the Archive Control File entry for each data set archived and if you have a tape management system it will honor the expiration specified. See Section 51.01 for details on retention of ABR Archive Backups.
DISP=(NEW,KEEP)	is required; do not specify CATLG since ABR handles cataloging of output files internally.
FREE=CLOSE	do not specify FREE=CLOSE since it will cause ABR to fail when it tries to create a second file on the tape.

```
EXAMPLE : //TAPE1      DD      DSN=ABR1,UNIT=3590,DISP=(NEW,KEEP),
           //          VOL=( , , 255 ),LABEL=EXPDT=99000
```

If multiple disk volumes are dumped to a given TAPEX DD, ABR will create multiple files on the tape (or tape aggregate if more than one tape volume is used), one file for each disk volume, using the naming convention in [Section 51.01](#). By default ABR will create as many as 255 files on a tape or aggregate before starting over with file 1 on a fresh scratch tape, but that limit can be overridden by the MAXFILE= operand, up to 4095 files. Larger MAXFILE= values can be used with high-capacity tapes such as IBM Magstar and StorageTek Redwood and 9840.

DCB parameters are not required and should be omitted. However, tape unit hardware compaction (sometimes called IDRC, available on most tape cartridge drives) can be requested by adding DCB=TRTCH=COMP to your DD statement; tape hardware compaction may be the default depending on local MVS options. For tapes attached by ESCON or FICON channels, Innovation recommends use of tape hardware compaction instead of FDR software compression (the COMPRESS= option).

51.04 CONTINUED . . .

The IBM Virtual Tape Server (VTS) and similar products from other vendors are supported. In a VTS, data written to "tape" is really written to disk internal to the VTS and is later moved to high-capacity tapes such as Magstars, resulting in much better physical tape utilization. When a tape is required for input, the data is staged back to the internal disk. VTS users may wish to specify MAXFILE=1 to avoid having to stage many tape files during a recall when only one file is actually needed, however, you may not want to use MAXFILE=1 if a second backup copy is directed to real tapes. MAXFILE=1 may require a larger number of virtual volumes in the VTS.

LAST TAPE Option:

The LAST TAPE option of ABR allows you to add backup files to a tape or tape aggregate created by a previous ABR step (even if that step is in another job and even if it was run on a previous day). This option is controlled entirely through JCL. To request LAST TAPE, the TAPEX DD is similar to that described above except that you specify:

DSN=FDRABR.LASTAPE.xxxxxxxx – the last index is 1 to 8 characters of your choice; you may have multiple LASTAPE files for various purposes. This name will be cataloged in the ABR catalog to record the tape volume serial and file number where ABR is to start its output on the next run.

DISP=(MOD,KEEP) – this tells ABR to locate the FDRABR.LASTAPE.xxxxxxxx file in the catalog, verify that the file exists on the output tape, and begin outputting to the tape at that point. If the name is not cataloged, ABR will call for a scratch tape and begin at file 1, so you can force ABR to start on a new tape simply by uncataloging the LASTAPE name. If you specify NEW instead of MOD, ABR will ignore the LASTAPE file and use scratch tapes (but it will still record the new LASTAPE file for future use).

VOL= – volume serials should not be specified, but the volume count should be given, e.g.,
VOL=(, , 255).

```
EXAMPLE : //TAPE1      DD      DSN=FDRABR.LASTAPE.ABR1,
                  //      UNIT=3490,DISP=(MOD,KEEP),
                  //      VOL=( , , 255 ), LABEL=EXPDT=99000
```

On cartridge drives, ABR will record the location of the LASTAPE file and will use hardware high-speed positioning when positioning to add files to the backup tape.

Disk Output:

You may request that ABR create Archive Backup files on disk. This is often used for COPY1, with a short retention such as 15 to 60 days so that recalls requested shortly after a data set is Archived can be completed without a tape mount (if you have an Automated Tape Library, recalls from tape may be almost as fast as from disk).

The TAPEX DD simply points to one or more disk volumes; ABR internally handles naming, allocating and cataloging the required backup files. [See Section 51.01](#) for the file names that ABR will create. Usually these are volumes which are dedicated to ABR use. **They cannot be SMS-managed volumes.**

```
EXAMPLE : //TAPE1      DD      UNIT=3390,DISP=OLD,
                  //      VOL=SER=(ARC001,ARC002,ARC003,ARC004,...)
```

ABR will start allocating backups on the first volser provided and will move down the list as the volumes fill up. By default, ABR allocates its backup files with:

```
SPACE=(TRK,(100,100),RLSE,ALX)
```


51.04 CONTINUED . . .

The ALX operand specifies that the backup data set will be allocated using the 5 largest extents on the volume that are at least 100 tracks large, so that it will have a minimum size of 100 tracks and a maximum of the size of the disk volume. RLSE requests that any unused space be released at the end of the backup. Since ABR can't predict how big the backup data set is going to be, this technique gets all or most of the free space on the target volume and releases it at the end. If the space allocated is still not sufficient, it will extend the backup file to additional disk volumes.

You may override the default allocation by specifying the SPACE= operand on the TAPEX DD. RLSE will be forced by ABR even if you do not specify it. You may want to specify ALX.

Do not specify a DSN= parm unless the POOLDISK option is used. You may specify RETPD= or EXPDT= to specify the expiration date of the backups; this is recorded in the Archive Control.

POOLDISK Option:

When outputting to disk, you may use the ABR POOLDISK option, allowing ABR to manage a pool of disks used for ABR backup data sets. Like LASTAPE, this is invoked through JCL options by specifying DSN=FDRA BR.POOLDISK.xxxxxxxx on the TAPEX DD, where xxxxxxxx is any name of your choice (allowing you to have multiple pools for various purposes). Otherwise the JCL is the same as just described for disk output.

```
EXAMPLE : //TAPE1      DD      DSN=FDRA BR.POOLDISK.ARCHIVE1,
           //              UNIT=3390,DISP=OLD,LABEL=RETPD=30,
           //              VOL=SER=(ARC001,ARC002,ARC003,ARC004,...)
```

Alternately you may catalog the POOLDISK data set with IDCAMS or IEFBR14, pointing to the volumes in the pool; then the volume list may be updated at any time without modifying the ABR JCL.

```
EXAMPLE : //TAPE1      DD      DSN=FDRA BR.POOLDISK.ARCHIVE1,
           //              DISP=OLD,LABEL=RETPD=30
```

In every ABR run, ABR will sort the POOLDISK volume list by the amount of free space available on each volume so that the volumes with the most available space will be used first (if the POOLSORT=NO option is specified ABR will use the volumes in the order specified, but will remember from run to run which volume it used last).

NOTE: When disk is used as a backup medium, a duplicate copy (TAPExx) to a real tape unit should be specified for safety. Retention of the backup on DISK should be for a short duration and the retention of the tape backup should be for as long as the archived data sets are to be kept. The retention period is coded on the TAPE DD statements.

WARNING: Tapes created by ABR cannot be copied using normal copy programs. Use the INNOVATION provided program, FDRTCOPY, to copy ABR tapes.

Simulation and Superscratch:

If Superscratch (DUMP TYPE=SCR), simulation of Superscratch (SIM TYPE=SCR) or simulation of Archive Backup (SIM TYPE=ARC) is requested, you must specify:

```
//TAPE1 DD DUMMY
```

51.04 CONTINUED . . .

TAPExx DD STATEMENTS Specifies that ABR is to create a duplicate backup (COPY 2). xx must specify the same character twice with xx corresponding to the x of the TAPEx statement.

EXAMPLE : //TAPE11 DD DSN=FDR11,DISP=(,KEEP),UNIT=TAPE
 will produce a tape which contains the same backup as

 //TAPE1 DD DSN=FDR1,DISP=(,KEEP),UNIT=TAPE

NOTE: TAPExx may specify any of the options as documented for TAPEx, including LASTAPE. However, a unique LASTAPE name must be used for each TAPE DD in the job step.

ABR will read the disk once and write the same data to TAPEx and TAPExx concurrently. These are known to ABR as COPY 1 (TAPEx) and COPY 2 (TAPExx). If TAPEx is on disk, TAPExx can be disk or tape. If TAPEx is tape, TAPExx must also be tape.

ABRARDQ DD STATEMENT Specifies the remote queue data set for Archive Backup. This DD statement is optional. If specified, ABR will read the control statements contained within, if any, and append these statements to the SYSIN data set. The SYSIN data set must contain at least a DUMP TYPE=ARC statement. After reading the control statements, ABR will reset the file to null (empty) except on SIM. DISP=SHR should be specified, since ABR internally serializes this data set.

CAUTION: This ABR execution must process all of the volumes which contain the data sets specified in the remote queue. This can easily be done by use of the ONLVOL operand on the DUMP statement.

SYSIN DD STATEMENT Specifies a data set containing the control statements for ABR. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY

.

51.05 ARCHIVE DUMP or SCRATCH STATEMENT

DUMP D SIM	TYPE=ARC SCR *	,MIGRAT=YES NO	GENERAL ARCHIVE SELECTION CRITERIA
	,ALTINDEX=NO	,MGMTCLAS=(mgmtclas,..) *	
	,ARCBACKUP=ABR DSF NO	,NOINIT *	,ADATE=yyddd yyyyddd *
	,ARCCAT=ALL NO NORMAL	,NOUNCAT	,ADAYS=dddd *
	,BUFNO=nn MAX	,ONLINE *	,CRDAYS=dddd **
		,ONLVOL *	,EXPIRED *
	,COMPRESS=ALL COPY1 COPY2	,POOLSORT=NO	,IFNOTCAT *
	,DATA=USED ALL	,PRINT=ABR *	,MAXGDG=nnn *
	,DCT=YES NO	,RECALL=YES NO	,REMOTE ***
	,DSNENQ=NONE TEST USE HAVE *	,RETPD=dddd	* If a data set meets any one of these GENERAL ARCHIVE SELECTION CRITERIA, it will be archived unless it matches a SELECT or EXCLUDE statement.
		,RETPD2=dddd	
	,DYNARC	,RTC=YES NO	** CRDAYS= is not a selection criteria. When used with other criteria it restricts selected data sets to those created at least dddd days ago.
	,ENQ=ON OFF RESERVE *	,SCRATCH[=IBM] *	
	,ENQERR=BYPASS PROCESS *	,SCRATCH=CAT *	*** REMOTE is not really a selection criterion. It enables Archive to select data sets flagged for archive by the FDRABRUT program.
	,EXPD=NONE *	,SELTERR=NO YES *	
	,VEXPD=NONE *		
	,FORMAT=NEW SPLIT	,SMSCOMMAND=NO YES	
	,ICFCORE=nnnnnn *	,SMSCONSTRUCT=YES NO *	
	,MAXBTRKS=nnnnn	,SMSEXPIRE=NO YES PRT ALL	
	,MAXCARDS=nnnnn *	,SMSMANAGE=NO YES*	
	,MAXDD=nnnnn *	,SMSMINRET=30 nnnn	
	,MAXERR=nnnnn	,SMSTHRESHOLD=YES NO *	
	,MAXFILE=nnnn	,THRESHOLD=NO HIGH LOW nnn *	
		,VOLSORT=YES NO	
		* only the marked operands can be used with TYPE=SCR	

DUMP STATEMENT The DUMP statement activates the Archive Backup (TYPE=ARC) or Superscratch (TYPE=SCR) of data sets. Only one DUMP statement is allowed per execution of ABR.

ARCHIVE BACKUP DUMP TYPE=ARC activates the ABR Archive Backup facility. Based on Archive criteria operands ABR will back up and scratch data sets from a disk volume. The volume must be enabled for Archive. These data sets will be recorded in the Archive Control File. At the completion of the Archive step, ABR will optionally backup the Archive Control File as the last file of the last backup tape used.

51.05 CONTINUED . . .

SUPER- SCRATCH	DUMP TYPE=SCR activates the ABR Superscratch facility. This statement accepts the same criteria as Archive Backup. ABR will scratch and uncatalog any data sets selected without taking a backup . These data sets will not be recorded in the Archive Control File or the ABR scratch catalog. The volume must be enabled for Superscratch.
SIM STATEMENT	If SIM is specified, ABR will perform the Archive or Superscratch function in a simulation mode. The TAPEX DD statement must specify DD DUMMY. ABR will not perform the actual Archive or Superscratch operation, only reporting the data sets which would be archived in the PRINT VTOC format. This allows various options to be tested without danger of archiving or scratching undesired data sets.
DATA SET SELECTION	If the ABR input contains one or more SELECT statements, any data sets matching a given SELECT will be controlled by that SELECT, which may or may not select it for Archive or Superscratch. If you supply DUMP statement operands from the box labeled "General Archive Selection Criteria", they will be used to make decisions about data sets not controlled by any SELECT statement. See Section 51.02 for details.
OPERANDS	<p>TYPE= Specifies the type of operation to be performed. It is required.</p> <p>ARC – an Archive Backup will be done. Data sets selected will be backed up, recorded in the Archive Control File, and scratched and uncataloged.</p> <p>SCR – A Superscratch will be done. Data sets selected will be scratched and uncataloged. No backup will be created.</p> <p>See Section 51.02 for details on the selection of data sets for Archive and Superscratch.</p> <p><i>WARNING: DUMP TYPE=SCR IS A VERY DANGEROUS FUNCTION SINCE THERE IS NO BACKUP OF THE DATA SETS. A SIMULATION SHOULD BE RUN PRIOR TO THE ACTUAL SUPERSCRATCH RUN.</i></p> <p>ADATE= Any data set which has not been opened since the date specified will be archived or scratched. The last reference date is stored in the Format 1 DSCB of a data set by IBM when it is opened. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).</p> <p>ADAYS= Any data set which has not been opened in dddd days will be archived or scratched. The last reference date is stored in the Format 1 DSCB of a data set by IBM when it is opened. The value specified will be subtracted from today's date to calculate a date dddd days in the past. If the result is greater than the last reference date, the data set will be selected.</p> <p>NOTE: ADAYS and ADATE are mutually exclusive. One of the above must be specified to activate archiving based on the last reference date.</p> <p>If you want to archive by last reference date, but only within a particular data set group, specify ADATE or ADAYS on a SELECT Statement, and do not specify ADATE or ADAYS on the DUMP Statement.</p>

51.05 CONTINUED . . .

ALTINDEX=	<p>NO – ABR will backup to the Archive Backup any ICF clusters that have an associated alternate index on the same volume as the primary, but will not delete them.</p> <p>Default is that ABR will not test for alternate indexes.</p>
ARCBACKUP=	<p>Specifies how ABR is to backup the Archive Control File at the completion of an Archive run.</p> <p>ABR – an ABR Volume Backup of just the ACF to be taken. This will increment the cycle and possibly the generation of the volume on which it resides (An full-volume backup might be forced).</p> <p>DSF – a DSF backup of the ACF is to be taken. ABR will create and catalog this backup using the name of the Archive Control File itself, except that the index level 'ARCHIVE' will be changed to 'ARCBKUP' for TAPEX and 'ARCBKU2' for TAPEXX. This cataloged name will only keep track of the most current backup of the Archive Control File. This option cannot be used if the Archive Control File has a non-standard name not containing an index level of 'ARCHIVE'. ARCBACKUP=DSF should not be used when archiving to disk since duplicate names will be created on the disk.</p> <p>NO – the Archive Control File will not be dumped at the end of the ARCHIVE.</p> <p>Default is ABR if the Archive is to tape; NO if the TAPEX backup is to DISK.</p>
ARCCAT=	<p>Specifies how ABR is to handle cataloging of backup files created by Archive Backup.</p> <p>ALL - all backup files will be cataloged in the ABR catalog.</p> <p>NO - no backup files will be cataloged.</p> <p>NORMAL - the first backup file created on each TAPEX or TAPEXX DD will be cataloged. Backups are also cataloged when the output volume list is different from that of the previous backup to the same DD, for example, when a tape fills up and a new scratch tape is mounted.</p> <p>Default is NORMAL, except that ARCCAT=ALL is forced if the backup is on disk or EXPDT=99000 (catalog control to some tape management systems) is specified on the TAPEX or TAPEXX DD.</p>
BUFNO=	<p>specifies how many buffers will be used for dumping each disk volume. Each buffer holds one disk track. The buffers acquired will be divided into 2 sets in order to overlap input and output I/O operations; each disk I/O will read disk tracks into one half of the buffers.</p> <p>MAX – buffers sufficient to read 1 cylinder of the input disk are acquired.</p> <p>nn - the specified number of buffers is acquired.</p> <p>Default: MAX. Innovation recommends that you do not override the default. However, BUFNO=2 will be forced when a backup (output) data set is on disk.</p>

51.05 CONTINUED . . .

COMPRESS=	<p>Controls the use of FDR software compression. Values for COMPRESS= are:</p> <p>ALL – the backup file for both copies (TAPEX and TAPEXX) is to be compressed.</p> <p>COPY1 – only the backup on TAPEX DD statements will be compressed.</p> <p>COPY2 – only the backup on TAPEXX DD statements will be compressed.</p> <p>See "Memory Requirements" in Section 50.03 for the additional storage required by COMPRESS=. COMPRESS= is ignored if DCT=YES is also specified.</p> <p>Default: backups will not be compressed.</p> <p>COMPRESS is recommended for backups to disk files, and for tape backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of IDRC (tape hardware compression) usually results in better performance.</p>
CRDAYS=	<p>ABR is not to select any data set based on the GENERAL ARCHIVE SELECTION CRITERIA on the DUMP statement unless it was created at least n days ago.</p> <p>Default is 2 days if IFNOTCAT was specified; otherwise default is 0 days.</p> <p>NOTE: CRDAYS on the DUMP statement functions as an EXCLUDE facility for the GENERAL ARCHIVE SELECTION. If none of the other criteria is specified (ex: ADAYS, IFNOTCAT) CRDAYS will be ignored. If you wish to select data sets solely based on creation date specify CRDAYS on SELECT statements.</p>
DATA=	<p>USED - only the used portion of PS (physical sequential) and PO (partitioned, PDS) data sets will be backed up. On most volumes, this will make the dump run faster.</p> <p>ALL - all allocated tracks of all selected data sets will be backed up.</p> <p>Default: USED.</p>
DCT=	<p>DCT= is valid only if you are licensed for FDR InstantBackup. It will be honored only if the disk being backed up is in a disk subsystem with the HSDM option (High Speed Data Mover). HSDM allows FDR to backup and restore the internal compressed images of disk tracks, improving backup elapsed times up to 60%. It can also be specified as DUMPCOMPRESSEDTRACK=.</p> <p>YES – use HSDM for any volume where the disk hardware has the HSDM feature installed. Normal backup will be used for other volumes.</p> <p>NO – do not use HSDM.</p> <p>Default is NO.</p> <p>Note that DCT=YES implies RTC=YES; see the description of RTC= for its benefits.</p>

51.05 CONTINUED . . .

DSNENQ=

Specifies whether all of the data sets selected for Archive or Superscratch will be ENQed. See "Data Set Enqueue Option" in [Section 51.03](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a warning message is issued for the data set and it will be bypassed. A successful ENQ will prevent any other task from attempting to use the data set until the backup of that volume is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be dumped. The options for DSNENQ= are:

USE – The data sets will be enqueued for the duration of the backup from this disk volume. This is the most frequently used option.

TEST – The data sets will only be tested to see if they are enqueued to another task at the time that the dump from this volume starts.

NONE – No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: NONE.

Recommendation: use DSNENQ=USE to be sure that the data sets being archived or scratched are not in use by some other task. You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements..

DYNARC

For Archive Backup, ABR is to dynamically allocate and use the common Archive Control File (ACF) whose name is stored in the FDR Global Option Table. The default name for this ACF is 'FDRABR.ARCHIVE', but your site may have changed it during FDR Installation. It will be dynamically allocated using DD name ARCHIVE#. Since all data sets which are archived for Auto-Recall (RECALL=YES) must be recorded in the common ACF, this option will be used for most Archive Backup jobs.

Default is that the user must supply an ARCHIVE DD statement pointing to the Archive Control File in which these backups are to be recorded.

51.05 CONTINUED . . .

ENQ=

Specifies whether an ENQ should be done on the VTOC of each disk volume while data sets from it are being processed. See “VTOC Enqueue Option” in [Section 51.03](#) for more details.

ON – the VTOC of each disk volume will be ENQed during its backup. This ENQ may be effective only on the system where the backup is executing; other systems may still be able to update the VTOC.

RESERVE – in addition to the ENQ, a hardware RESERVE will be issued on each disk volume during its backup. This is meaningful only on a system with “shared DASD” where the disks can be accessed by another MVS system. On the system where FDR is executing, an ENQ for (SYSVTOC,volser) is done, but other systems will be unable to read or write any data on the volume.

OFF – the VTOC will not be enqueued or reserved during the backup.

Default is ON.

ENQERR=

Specifies processing if the DSNENQ= option finds that a data set is in use (enqueued):

BYPASS – do not archive or scratch an active data set.

PROCESS – archive or scratch a data set even if it is active (a warning message will still be produced). If IBM SCRATCH or DELETE is used to scratch the data set (see SCRATCH=), the scratch may fail anyway if the data set is enqueued by another task, so **ENQERR=PROCESS is not recommended**.

Default: BYPASS.

EXPD=
VEXPD=

These control what ABR does when it finds that a data set about to be scratched has a non-zero expiration date which has not yet been reached (unexpired). It is not used for selection of data sets (see the EXPIRED operand).

EXPD=NONE – ABR will override the expiration date for all data sets and scratch them.

VEXPD=NONE – ABR will override the expiration date for ICF VSAM clusters and delete them. Either EXPD=NONE or VEXPD=NONE must be specified for ABR to delete unexpired ICF VSAM clusters.

Default is that ABR will not scratch unexpired data sets. For unexpired non-VSAM data sets, ABR will ask the operator for permission before scratching the data set (console message FDRW23). For unexpired VSAM files, ABR will not issue an operator message and the scratch will fail.

EXPIRED

ABR is to archive or scratch data sets if their expiration date is less than or equal to today's date. ABR will not select data sets with an all-zero expiration date or a special expiration (99.000, 99.365, 99.366) unless they are selected by other criteria. The expiration date is not checked on VSAM files unless they are SMS-managed.

51.05 CONTINUED . . .

FORMAT=

Specifies the format of the sequential backup file.

NEW – the backup will be created using a maximum of a 56K blocksize. A block will contain the image of one or more tracks from the input disk.

SPLIT – the backup will be created using a maximum blocksize of 32K. For blocks that would be more than 32K, they are written as 2 blocks of 32K or less. However, **FORMAT=SPLIT** causes use of a disk I/O technique which is less efficient than that used by **FORMAT=NEW**, which will impact backup performance.

WARNING: *If you use a normal copy program (ex: IEBGENER) to copy a backup file created with **FORMAT=NEW**, you will not get any error messages, but the resulting tape will not be usable for a restore. Tapes in the new format must only be copied with the utilities **FDRTCOPY** and **FDRTSEL** (See Section 60).*

Default: **NEW** if backup on tape – **SPLIT** if backup is on disk.

ICFCORE=

Specifies that the size of the table used to store the ICF VSAM cluster and component names. The backup must save all of the component names and their associated clusters which exist on the current input disk volume, in order to match up VTOC DSCBs (with the component name) to cluster names for selection. nnnnnn is specified in bytes and must be large enough to contain all the VSAM names (depending on the length and number of names).

NOTE: Specifying **ICFCORE=** will increase the backup region requirement by the value specified. The default value uses one of the dump buffers and imposes no additional memory requirement.

Default: 53248, which will hold about 650 components.

IFNOTCAT

ABR is to archive non-VSAM data sets which are either not cataloged at all, or cataloged to a volume other than the one they are selected from. If the proper catalog is not available or any other catalog error exists, **ABR** will not use this selection criterion. By default, **ABR** will not archive uncataloged data sets if they are less than 2 days old, unless overridden by using the **CRDAYS=** operand. A time delay is necessary to avoid archiving active data sets which will be cataloged at the end of the step or job that is creating them.

NOTE: **IFNOTCAT** can be very expensive to run if a large number of data sets must be checked by **ABR**, due to the large number of expensive catalog accesses.

Default is that the catalog status of data sets is not checked.

MAXBTRKS=

If more than the specified number of disk tracks are selected for Archive Backup from a given disk volume, **ABR** will make multiple passes on that disk, creating multiple Archive Backup files. This reduces the size of individual archive backups, and will improve recall performance. The size of each backup will vary, since archived data sets will not be split among backup files. **MAXBTRKS=65536** or more will disable this feature.

Default is 4096 tracks.

51.05 CONTINUED . . .

- MAXCARDS=** Enables ABR to accept additional SELECT, EXCLUDE and MOUNT statements during this execution. You can specify values up to 65535, but values over 100 will increase the region required by ABR.
- Default is 100 statements.
- MAXDD=** Specifies the maximum number of disk volume serials that ABR can process in this execution. ABR builds its volume list from DISKxxxx DD statements and/or ONLINE volumes if ONLINE, ONLVOL or MOUNT is coded. Once this limit is reached additional volumes will be ignored.
- Default is 256 volumes, unless overridden by the MAXONLINE option in the FDR Global Option Table ([see Section 90](#)).
- MAXERR=** The number of tape or disk errors that are permitted prior to abending the operation. MAXERR may specify a value from 1 to 9999 errors. Each error will be indicated by a message and possible MINI DUMP.
- Default: 20 errors.
- WARNING: On ARCHIVE operations, although MAXERR will allow the DUMP operation to proceed without ABENDING, no data sets will be scratched or uncataloged from a given disk volume if any tape or disk errors occur during its backup. This option should only be used when necessary and with care.***
- Default is 20 errors.
- MAXFILE=** Specifies the maximum number of files ABR will create on tape. May specify from 1 to 4095. When the maximum file number is exceeded, ABR will start a new scratch tape using file sequence number 1. A MAXFILE= over 255 may be appropriate when you are outputting to high-capacity tapes such as IBM 3590 Magstar or StorageTek Redwood and 9840.
- Default is 255 unless overridden in the FDR Global Option table (See [Section 90](#)).
- If you are outputting to an IBM Virtual Tape Server (VTS) and a similar product from another vendor, you may wish to specify MAXFILE=1 to avoid having to stage many tape files when only one file is actually needed. [See Section 51.04](#).
- MAXGDG=** ABR is to Archive any GDG (Generation Data Group) data sets which exist on the disk volumes to be scanned and are not within n generations of the most current generation recorded in the catalog. n can range from 1 to 255. (Example: MAXGDG=2 causes ABR will keep the two most current generations of every GDG, archiving all other generations if found). It will also archive any generations which have **higher** generation numbers than the current generation in the catalog.
- WARNING: If you use MAXGDG, you should specify CRDAYS=2 or higher, or specify DSNENQ=, to avoid archiving GDG generations that are in the process of being created.***

51.05 CONTINUED . . .

MGMTCLAS=	<p>On SMS-managed volumes, only those data sets whose SMS management class matches one or the class names specified will be considered for selection. Multiple class names are specified in parenthesis. You may also specify multiple MGMTCLAS= parameters with the same result. SMSMANAGE=YES must also be specified.</p>
MIGRAT=	<p>YES – specifies that when ABR recatalogs an archived data set for Auto-Recall (see the RECALL operand), the first volser in the catalog entry will be changed to "MIGRAT".</p> <p>NO – specifies that the recatalog for Auto-Recall will preserve the original volume serial in the catalog entry.</p> <p>Default is NO unless overridden in the FDR Global Option Table (See Section 90).</p> <p>Note: MIGRAT=YES, either specified or from the Global Option Table, is required if SMS-managed data sets are being archived. It is strongly recommended for all Archive jobs.</p>
NOINIT	<p>For SIM only. ABR is to perform the simulation on the volumes specified, even if they are not initialized for ABR processing or disabled for Archive or Superscratch processing. An ABR error message will still be issued but the volumes will be processed.</p> <p>Default: simulation cannot be performed on a volume unless contains an ABR Model DSCB (see Section 50.01) that is enabled for Archive or Superscratch.</p>
NOUNCAT	<p>Specifies that ABR is not to uncatalog any non-VSAM data sets that are archived or scratched.</p> <p>Default: ABR will uncatalog all non-VSAM data sets archived or scratched unless they are cataloged to a different volume. For multi-volume data sets, the catalog entry is not deleted until the data set has been archived or scratch from each of its volumes. If RECALL=YES is in effect for Archive, ABR will update the catalog entry.</p> <p>ICF VSAM clusters are always uncataloged when they are deleted.</p> <p>NOTE: NOUNCAT and RECALL=YES are mutually exclusive.</p>

51.05 CONTINUED . . .

**ONLINE
ONLYVOL**

If these operands are omitted, ABR will only process volumes specified by DISKxxxx DD statements or ABR MOUNT statements.

ONLINE – specifies that ABR is to process every volume that is ONLINE to the system; if there are more than 255 online volumes, you will need to specify MAXDD= to increase that limit. With ONLINE, there is usually no need to specify DISKxxxx DD statements or MOUNT statements. However, if you have certain disk volumes which you want processed first, specify them on DISKxxxx DD statements; other online volumes will be appended to the list created from the DD statements.

ONLYVOL – specifies that ABR is to scan all SELECT statements for the operands VOL=, VOLG=, or CATDSN= (for volume selected from the catalog). These volumes, if online, will be appended to the ABR volume list.

If ABR finds that the user specified EXCLUDE statements with ALLDSN and VOL or VOLG, these volumes will not be included in the ONLINE list unless a DISKxxxx DD statement is present for this volume. This is a way of excluding certain volumes, for example, EXCLUDE

ALLDSN, VOLG=WORK.

POOLSORT=

NO – if the POOLDISK option is specified on one or more TAPEX or TAPEX DDs in the ABR step (see [Section 51.04](#)) the volumes in the POOLDISK list will be used in the order specified in the JCL.

Default: the POOLDISK volumes will be sorted based on the amount of available space on each, first using those with the most free space.

PRINT=

ABR – Specifies that ABR is to use the PRINT VTOC format of program FDRABRP (See [Section 53](#)) for each data set that was selected.

Default: the names of selected data sets are simply listed on SYSPRINT. However, SIM (simulation) defaults to PRINT=ABR.

RECALL=

YES (or just **RECALL**) – ABR is to keep archived data sets cataloged for Auto-Recall by the ABR Catalog Locate Exit (see "Auto-Recall" later in [Section 51](#)), but only if the data set is currently cataloged to the volume it was backed up from. If the data set is cataloged to another volume or not already cataloged, it will not be cataloged for RECALL. If an ICF VSAM cluster is archived, ABR will recatalog the cluster name as non-VSAM. Unless you have not installed the ABR Locate Exit or have special reasons for archiving data sets that cannot be automatically recalled, RECALL=YES is recommended.

NO – specifies that ABR will uncatalog Archived data sets if they are cataloged to this volume unless NOUNCAT is specified.

Default is NO unless overridden in the FDR Global Option table (See [Section 90](#)).

Recommendation: if you plan to use RECALL=YES for most of your archived data sets, set the default to YES in the FDR Global Option Table and use RECALL=NO for those few jobs where it is not desired.

NOTE: RECALL=YES and NOUNCAT are mutually exclusive.

51.05 CONTINUED . . .

REMOTE

ABR is to archive the data sets requested by the remote queue (program FDRABRUT). REMOTE is ignored on TYPE=SCR. REMOTE is required only if data sets were flagged on disk with the DISKUPDATE=YES option of FDRABRUT, and is automatically included if any other GENERAL ARCHIVE SELECTION CRITERIA keywords are specified on the DUMP statement.

**RETPD=
RETPD2=**

RETPD= specifies the number of days (1 to 9999) that COPY1 (TAPE_x) backups will be kept. ABR will calculate an expiration date from this value. If RETPD2= is omitted, this same expiration will apply to all COPY2 backups created in the same step.

RETPD2= specifies the number of days (1 to 9999) that COPY2 (TAPE_{xx}) backups will be kept. ABR will calculate an expiration date from this value.

[See Section 51.01](#) for a discussion of backup retention. Also see SMSEXPIRE= in this section.

Default is 365 days unless overridden by RETPD= or EXPDT= on individual TAPE DD statements.

RCT=

It can also be specified as READTRACKCCW=.

YES – use READ TRACK CCWs to read disk data tracks. RTC=YES also causes:

- up to 1 cylinder of disk data is read at a time.
- FDR buffers are moved above the 16MB line (about 2MB per concurrent backup), allowing more concurrent backups to be run in one step.
- the elapsed time of ABR backups when the backup data set is itself on disk is significantly improved.

NO – use other CCWs to read disk data tracks.

Default is NO.

Innovation recommends the use of RTC=YES (or, if your disk subsystem supports the HSDM (High Speed Data Mover) hardware option, DCT=YES) to improve backup performance when creating backups on disk. DCT=YES implies RTC=YES.

51.05 CONTINUED . . .

SCRATCH

SCRATCH or **SCRATCH=IBM** – specifies that ABR will issue the IBM SCRATCH SVC or VSAM DELETE to scratch each data set.

SCRATCH=CAT – same as SCRATCH except that ABR is to record this data set in the ABR scratch catalog (if enabled) if it has a current ABR Volume Backup ([see Section 50](#)), in addition to following all of the above rules. This may be useful with Superscratch to allow ABR to restore a backup copy of the data set as long as the backup is still available.

Default is to use the IBM SCRATCH SVC or VSAM DELETE if the disk volume has an active INDEXED VTOC (VTOCIX) or if ICF VSAM clusters are to be deleted, otherwise ABR will use its own technique for scratching data set entries from the VTOC. This technique involves zeroing the DSCBs of the scratched data sets; at the end of successful Archive or Superscratch, ABR will force the IBM VTOC conversion routine to recalculate the new free space by setting the DOS bit in the Format 4 DSCB and allocating a dummy data set.

NOTE: The dummy data set that ABR allocates to invoke the VTOC conversion routine has a name that starts with FDRABR.V. If your installation has a data security system, then the user running the Archive or Superscratch job must be authorized to create this dsname. If this is not possible then specify SCRATCH=IBM.

SELTERR=

Specifies what will happen at step termination if ABR finds that a SELECT or EXCLUDE statement was never referenced (no data set on any input disk was selected by the statement):

NO – a condition code or ABEND is **not** to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.

YES – a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.

Default: YES unless overridden in the FDR Global Option Table ([See Section 90](#)).

SMSCOMMAND=

NO – when selecting data sets from SMS-managed volumes by management class attributes (SMSMANAGE=YES), SELECT/EXCLUDE statements are NOT to be honored. **Data sets will be selected based ONLY on management class attributes.** Data sets with the attribute COMMAND-OR-AUTO-MIGRATE=COMMAND will not be selected.

YES – for SMS data sets with the attribute COMMAND-OR-AUTO-MIGRATE=COMMAND or BOTH, SELECT/EXCLUDE statements in the ABR run may be honored ([See Section 70](#) for details). Note that SMSCOMMAND=YES should be specified on an ARCHIVE run that takes input from the ARCHIVE REMOTE QUEUE if user requests for SMS-managed data sets are to be honored.

SMSCOMMAND= applies only to DUMP TYPE=ARC.

Default is NO.

51.05 CONTINUED . . .

SMSEXPIRE=

NO – specifies that no special processing be done for the expiration of Archived SMS-managed data sets. They receive the normal ABR expirations set from JCL parameters or the RETPD= operand.

YES – specifies that Archived SMS-managed data sets selected based on their management class attributes will receive individual expirations calculated from the attributes of their associated SMS management class. Honored only if SMSMANAGE=YES is also specified. See [Section 70](#) for more details.

PRT – same as YES, except that the calculated expirations for each Archived SMS-managed data set will be displayed.

ALL – same as YES, except that data sets which were not selected because of management class attributes but which were selected because of ABR SELECT statements will also receive individual expirations calculated from the attributes of their associated SMS management class.

SMSCOMMAND=YES must also be specified for the SELECT statements to be honored. If you want to print the assigned expirations, specify SMSEXPIRE=ALL, SMSEXPIRE=PRT.

SMSEXPIRE= applies only to DUMP TYPE=ARC.

Default is NO.

SMSCONSTRUCT=

YES – ABR is not to process SMS-managed volumes unless their associated SMS storage group has the attribute AUTO-MIGRATE=YES (and the ABR Model DSCB for the volume is also enabled for ARCHIVE or SUPERSCRATCH as appropriate).

NO – ABR will bypass selected volumes based only on the ARCHIVE or SUPERSCRATCH options in the ABR Model DSCB, ignoring the SMS storage class attribute.

Default is YES.

SMSMANAGE=

NO – specifies that ABR is not to select data sets from SMS-managed volumes based on the attributes of their SMS management class. ABR Archive Selection Criteria will be used to select both SMS and non-SMS data sets.

YES – specifies that data sets on SMS-managed volumes are to be selected for ARCHIVE or SUPERSCRATCH based primarily on the attributes of their associated SMS management class ([See Section 70](#) for complete details). Non-SMS data sets processed in the same run will use ABR Archive Selection Criteria.

Default is NO.

WARNING: if you specify SMSMANAGE=YES, the archive PROTECT List is ignored, except for entries with ALLDSN (such as PROTECT ALLDSN, VOLG=PROD) which cause the volume to be completely bypassed.

51.05 CONTINUED . . .

- SMSMINRET=** If SMSEXPIRE= is also specified, this specifies the minimum COPY2 retention to be calculated. This insures a reasonable retention for data sets which would otherwise expire very quickly under management class rules.
- Default is 30 days.
- SMSTHRESHOLD=** **YES** – ABR will use threshold values from the SMS storage group for SMS-managed volumes.
- NO** – thresholds from the ABR Model DSCB will be used for thresholding on all volumes.
- SMSTHRESHOLD= is valid only if THRESHOLD=HIGH/LOW is specified.
- Default is YES.
- THRESHOLD=** Specifies whether ABR will bypass volumes because their current allocation (percentage of tracks allocated to data sets) is below a certain threshold. For non-SMS volumes the thresholds may be stored in the ABR Model DSCB on the volume. For SMS-managed volumes, the threshold may be in the SMS storage group or may be taken from the ABR Model DSCB (see SMSTHRESHOLD= above).
- NO** – volumes are not to be bypassed because of allocation thresholds.
- HIGH** – the threshold used is the high threshold in the ABR Model DSCB or the SMS storage group.
- LOW** – the low threshold in the ABR Model DSCB or SMS storage group will be used.
- nnn** – is a number from 0 to 100 which will be used as the threshold (values from the ABR Model DSCB and SMS will not be used).
- Default is NO.
- Note:** the MAXBTRKS= operand may affected threshold processing. If a large amount of data is being archived from a volume, MAXBTRKS= causes ABR to backup only a fraction of it at a time, making several passes on the volume. The threshold will be evaluated at the start of each pass, so ABR may not select all of the eligible data sets. To avoid this, specify a large value for MAXBTRKS=. If you want this to occur, specify a smaller MAXBTRKS=

51.05 CONTINUED . . .

VOLSORT=

YES — if you have more than one TAPEX DD statement in this ABR step, ABR will sort the volumes to be backed up by the last digit of their MVS device address, to attempt to balance channel and control unit utilization during concurrent backups. In other words, all volumes whose MVS address is xxx0 will be processed first, then xxx1, etc. This only applies to volumes selected by the ONLINE or ONLVOL operands or MOUNT statements.

NO — volumes will be processed in the order that they are found during a scan of the system UCBs, which is usually (but not always) in MVS device address order.

Default is YES except that NO is forced if you have only one TAPEX DD statement.

The above does not apply to volumes selected by DISKxxxx DD statements which will always be processed first, in the order the DD statements appear in the JCL.

51.06 SELECT, EXCLUDE AND MOUNT STATEMENTS

SELECT S	DATA SET SELECTION CRITERIA DSN=filter * CATDSN=filter * DD=ddname * ABRBKUP=volser * TEMP * ALLDSN * ,CATALOG=catname ,MCATALOG=catname ,CATLIMITGDG=n ,CRDAYS=dddd * ,DATA=ALL ,DSORG=(xx,xx..) * ,GDG * ,PRTALIAS ,SIZE=nnnnn * ,UPDATE ,VOL=vvvvvv * ,VOLG=vvvvv * ,ALLVOL * * Data set must meet all of the marked criteria, if specified, to be selected by this statement. On an EXCLUDE statement, only these marked operands can be used.	ARCHIVE SELECTION CRITERIA ,ADATE=yyddd yyyyddd ,ADAYS=dddd ,EXPIRED ,IFNOTCAT ,MAXGDG=nnn If any of the above specified ARCHIVE criteria is met, the selected data set will be archived. If none are specified, the data set is unconditionally archived.
EXCLUDE X		

MOUNT **VOL=vvvvv|VOLG=vvvvv|STORGRP=storagegroup**

51.06 CONTINUED . . .

**SELECT/
EXCLUDE
STATEMENTS**

These statements select additional data sets to be archived or scratched or exclude them from processing.

The SELECT statement has 2 purposes with Archive and Superscratch.

1. it can be used to unconditionally select data sets to be archived or scratched. The SELECT statement will contain only Data set Selection Criteria operands from the first box on the previous page. As ABR scans the VTOC of every selected volume, it will compare each data set to those specifications and archive or scratch any data set which matches **all** of the criteria specified.
2. It can be used to specify the criteria for archiving or scratching special groups of data sets. In addition to the Data set Selection Criteria from the first box, it will contain one or more Archive Selection Criteria from the second box on the previous page. For every data set selected by the Data set Selection Criteria, ABR will archive or scratch only those data sets that match **one or more** of the Archive Selection Criteria.

An **EXCLUDE** statement identifies data sets from within those selected by SELECT statements which are not to be processed; it can also be used to exclude data sets which would otherwise be selected by Archive Selection Criteria on the DUMP statement.

All data sets in the VTOCs of DASD volumes processed by ABR will be compared to these statements to identify those to be processed; each data set will be compared to each control statement until a match is found. It must match **all** Data set Selection Criteria specified on the statement to qualify, e.g, if DSN= and DSORG= are both specified, it must be a data set with the right name, having the indicated organization. A maximum of 100 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since ABR will only archive or scratch data sets which are selected, EXCLUDE statements can be used to exclude certain data sets from within a larger group on a SELECT statement

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A . B . **
SELECT  DSN=A . **
```

Example 2. Select all data sets except partitioned (PDSs):

```
EXCLUDE ALLDSN , DSORG=PO
SELECT  ALLDSN
```

You can also use an EXCLUDE ALLDSN statement with VOL= or VOLG= to specify entire volumes which are not to be processed by ABR. This is a way to exclude certain volumes that would otherwise be selected by the ONLINE or ONLVOL operands or MOUNT statements.

[See Section 51.02](#) for more details on data set selection.

**MOUNT
STATEMENTS**

The **MOUNT** statement is used to specify additional volumes to be processed by ABR. VOL=, VOLG=, or STORGRP= must be specified. STORGRP= can be used only on systems with SMS active; it will select all volumes in the specified SMS storage group. MOUNT statements should follow SELECT/EXCLUDE statements, if any.

**ICF VSAM
CLUSTERS**

ICF VSAM clusters can be selected by specifying the fully-qualified base cluster name or matching on the base cluster with generic data set selection. When selected, all components of that cluster that exist on the volumes being processed will be archived or scratched, including alternate indexes and key range components. ABR will not examine ICF VSAM component names when processing SELECT/EXCLUDE statements; components will be selected only if their cluster name is selected. For further information, [see Section 80.13](#) "VSAM Special Considerations".

If a cluster spans multiple disk volumes, it will be properly archived or scratched only if **all** of the volumes containing a piece of the cluster are processed in this ABR step.

51.06 CONTINUED . . .

OPERANDS DSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the VTOCs of selected volumes.

Examples : DSN=USER1.JCL.CNTL
 DSN=**LIST
 DSN=PROD++.**.LIB*

DSN= cannot select GDGs by relative generation number; use DD= or CATDSN= for that purpose.

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= operand (you should specify the ONLVOL operand so that those volumes are automatically added to the ABR volume list for processing). Specification of a relative generation number for GDG data sets is supported, e.g.,
 CATDSN=A.B(-1)

If a filter is specified, then catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT CATDSN=dsname was present for each of them. It may be necessary to specify MAXCARDS=nnnnn if a large number of data sets are selected by the filter.

Additional considerations for CATDSN=filter are explained in [Section 80.14](#).

CATDSN= is supported only on SELECT statements. However, a preceding EXCLUDE statement with DSN= and/or VOL= can exclude data sets from selection by CATDSN=.

If the VOL= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

Examples : CATDSN=USER1.JCL.CNTL
 CATDSN=**MASTER(0)
 CATDSN=PROD++.**.LIB*

Normally CATDSN= will not display the data sets it selects from the catalogs, you will see the names only when ABR actually finds and selects the data sets in the VTOCs of the volumes they are cataloged to. To display all of the data sets selected specify PCATDSN=filter.

WARNING: depending on the filter specified, CATDSN= may need to search many catalogs.

51.06 CONTINUED . . .

DD= Specifies that a data set name is to be taken from a DD statement. The value is the DDNAME of a JCL statement in the ABR step. Using this option enables you to specify a non-standard data set name or a generation data set (GDG) relative generation.

Example: SELECT DD=DD1
 //DD1 DD DSN=A.B.C(0),DISP=SHR

Note that although DD= copies the data set name from the DD statement, it does not use the volume pointed to by that DD unless the DDNAME is DISKxxxx. The data set will be searched for only on the volumes processed by ABR; the volume containing the data set must be among them.

ABRBKUP Specifies that ABR is to select ABR backup data sets currently on disk which have expired (reached their expiration date). This operand can be used with Superscratch (TYPE=SCR) to clean up volumes containing ABR backups which are no longer needed.

If **ABRBKUP** is specified alone, then all expired ABR backup files will be selected.

If **ABRBKUP=volser** is specified then only backups containing data sets from volume "volser" will be eligible. "volser" may contain filter characters, such as ABRBKUP=TSO+++ to select volumes called TSO plus 3 digits ([see Section 80.14](#)).

NOTE: You cannot use SELECT ABRBKUP to move disk archive files to tape; archive backup files cannot themselves be archived. See FDRTSEL in [Section 60](#) to do this job.

TEMP Specifies that ABR is to select all non-VSAM data sets which have a name in the format which the system uses for temporary data sets. This may be used with TYPE=SCR to clean up temporary data sets which were not automatically deleted at job termination. It is recommended that CRDAYS=2 also be specified to insure that active temporary data sets are not selected.

NOTE: Neither SELECT ALLDSN nor General Archive Criteria on the DUMP statement will select temporary data sets; a SELECT TEMP statement is required.

ALLDSN Specifies that all the data sets encountered on the volumes specified are to be selected.

Note: DSN=** is equivalent to ALLDSN, except that, unlike ALLDSN, DSN=** will not exclude certain types of data sets as described in [Section 51.02](#).

NOTE: DSN=, CATDSN=, DD=, TEMP, ABRBKUP and ALLDSN are mutually exclusive. One and only one of them must be specified on each SELECT/EXCLUDE/PROFILE/PROTECT Statement.

ADATE= Any data set which has not been opened since the date specified will be archived or scratched. The last reference date is stored in the Format 1 DSCB of a data set by IBM when it is opened. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

51.06 CONTINUED . . .

- ADAYS=** Any data set which has not been opened in dddd days will be archived or scratched. The last reference date is stored in the Format 1 DSCB of a data set by IBM when it is opened. The value specified will be subtracted from today's date to calculate a date dddd days in the past. If the result is greater than the last referenced date, the data set will be selected.
- NOTE:** ADAYS and ADATE are mutually exclusive. One of the above must be specified to activate archiving based on the last reference date.
- CATALOG=**
MCATALOG= Specifies the name of a user catalog (CATALOG=) or alternate master catalog (MCATALOG=) to search when CATDSN= is specified. [See Section 80.14](#) for details.
- Default is that the catalog search will start with the active master catalog. User catalogs will be searched if their assigned aliases match the CATDSN=filter.
- CATLIMITGDG=** May be used with CATDSN=filter to limit the selection of GDGs from the catalogs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:
- n** will cause only the most recently created "n" generations to be selected.
- n** will cause only generation (-n) to be selected.
- Default is that all the generations of selected GDGs will be selected unless a relative generation number is specified at the end of the filter, e.g., CATDSN=filter(-2).
- CRDAYS=** data sets will not be selected unless they is at least dddd days old. A creation date is stored in the Format 1 DSCB of every data set when it is created. The value specified will be subtracted from today's date to calculate a date dddd days in the past. If the result is greater than the creation date, the data set will be selected.
- Default is 0 unless IFNOTCAT is specified then 2 days is assumed.
- DATA=** **ALL** – specifies that Archive will backup the entire allocated space of the selected data sets. Normally it will only process up to the last block pointer (end-of-file) on input PS or PO data sets. It can be used if the last block pointer of certain data sets is invalid.
- DSORG=** Specifies that these data sets is not to be selected unless DSORG matches one of the DSORGs specified. If more than one DSORG is specified, they must be enclosed in parenthesis.
- VALID DSORGs are:
- | | | | |
|-----------|------------------|----------------|----------------|
| DA – BDAM | PS – SEQUENTIAL | AM – ALL VSAM | EF – ICF VSAM |
| IS – ISAM | PO – PARTITIONED | UN – UNDEFINED | UM – UNMOVABLE |
- NOTE:** It is not recommended that ISAM or non-ICF VSAM files be selected for archive.

51.06 CONTINUED . . .

EXPIRED Specifies that ABR is to archive or scratch the data sets specified if the expiration date stored in their Format 1 DSCB is equal to or less than today's date. This option does not select data sets with an all-zero expiration date or a special expiration date (99.000, 99.365, 99.366). The expiration date is not checked on VSAM files, unless they are SMS-managed. If SMSMANAGE=YES is specified on a TYPE=SCR run then expired data sets on SMS-managed volumes will not be scratched unless EXPIRED is also specified.

GDG data sets will not be selected by this SELECT unless they are generations of a generation data group (GDG).

IFNOTCAT ABR is to archive non-VSAM data sets which are either not cataloged, or cataloged but to a volume other than the one they are selected from. If the proper catalog is not available or any other catalog error exists, ABR will not use this selection criterion. By default, ABR will not archive uncataloged data sets if they are less than 2 days old, unless overridden by using the CRDAYS= operand. A time delay is necessary to avoid archiving active data sets which will be cataloged at the end of the step or job that is creating them.

NOTE: IFNOTCAT can be very expensive to run if the system contains a large number of data sets, due to the large number of catalog accesses.

Default is that the catalog status of data sets is not checked.

MAXGDG= Specifies that ABR is to archive or scratch the data sets specified if they are generations of a GDG (Generation Data Group) and are not within n generations of the most current generation recorded in the catalog. n can be from 1 to 255. (Ex: MAXGDG=2, ABR will keep the two most current entries, archiving all others if found.) If you use MAXGDG=, you should also specify CRDAYS= or DSNENQ= to avoid selecting data sets that are in the process of being created.

PRTALIAS When used on a SELECT statement with CATDSN= will display all of the alias names and user catalogs that were searched.

SIZE= Data sets are not to be selected unless their size in tracks is equal to or larger than the size specified. For ICF VSAM clusters the size is checked for each of the components. If SIZE is to be specified in more than one SELECT statement for the same group name, they must be coded in descending sequence (the largest first).

STORGRP= Can only be used on the MOUNT statement, and only on systems with SMS (System Managed Storage) active. Will select all online volumes in the specified SMS storage group.

EX: MOUNT STORGRP=DBLARGE

will cause ABR to process all volumes in that storage group.

UPDATE Data set which match this SELECT statement will be selected only if the update indicator is set in their Format 1 DSCB, indicating that it has been updated since the last ABR Volume Backup ([Section 50](#)).

NOTE: an ARCHIVE job will NOT reset the update indicator; it will be reset only by an ABR Volume Backup.

51.06 CONTINUED . . .

VOL= Specifies a volume serial number (up to 6 characters) to which this statement applies. If not specified, ABR will examine all volumes processed for the data sets named on this statement. For MOUNT statements, it will add the named volume to the ABR volume list if it is online.

VOLG= Specifies a volume serial prefix (1-5 characters) to which this statement applies. If not specified, ABR will scan all volumes processed for the data sets named on this statement. For MOUNT statements, it will process the all online volumes which start with the prefix.

ALLVOL Specifies that the data sets indicated on the control statement apply to all the volumes encountered within this execution of ABR.

NOTE: VOL, VOLG, STORGRP and ALLVOL are mutually exclusive.

Default is ALLVOL.

51.07 ARCHIVE RESTORE JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform manual restore from Archive Backups.

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRABR), region requirement (REGION=), and optional PARM= operand. The minimum region required is 512K. However, some restore options, especially logical restore, may increase the region requirement, so a value of 1M or 2M is recommended; REGION=0M can be specified to get the largest below-the-line region available.

If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid RESTORE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=ARC'
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=ARC/ SELECT DSN=A.B.C'
```

If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

STEPCAT or JOBCAT DD STATEMENT For non-VSAM data sets, when FDR must catalog non-VSAM data sets, they will be cataloged in that user catalog instead of in the system catalog with the matching alias. This might be useful when you are creating a test system, to catalog test copies of production data sets in a test catalog.

However, there is one important exception: if the data set being cataloged is a GDG generation, the STEPCAT/JOBCAT must contain a GDG base for that GDG; if not, it will ignore the STEPCAT/JOBCAT and catalog into the regular aliased catalog, possibly deleting other valid generations.

For VSAM clusters, the target catalog is controlled by the ICFCAT= operand, described in [Section 51.08](#). With the proper ICFCAT= option, the STEPCAT/JOBCAT may be honored.

STEPCAT/JOBCAT should not be used if any data sets being restored are SMS-managed.

SYSPRINT DD STATEMENT Specifies the output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=120; the blocksize will default to 1210 on disk or tape.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

ARCHIVE DD STATEMENT Specifies the ABR Archive Control File. This DD is required for restores from Archive Backups unless the DYNARC operand is specified on the RESTORE statement. The standard name for this data set is FDRABR.ARCHIVE, but your installation may have chosen a different name. Specify DISP=SHR since ABR internally serializes access.

51.07 CONTINUED . . .

**TAPEx DD
STATEMENT**

Allocates an input tape drive to be used for the RESTORE operation. "x" may be 1 alphabetic (A-Z), numeric (0-9) or national (@ \$ in the US) characters. The national character # and its international equivalents are reserved for ABR dynamic allocation. This tape drive must be capable of reading all the tapes needed for this restore; if the tapes are in varying formats, such as 3480 and 3490E, the restore will fail (see "DYNTAPE note" below).

Although this DD allocates a tape drive, ABR will internally fill in the data set name, volume serials, and file number before it OPENS each tape file, so you do not have to provide any information about the backups to be read. However, MVS requires that you provide:

DSN= - provide any dummy data set name. It will not be used by ABR, but MVS will ENQ on this name so multiple ABR restore jobs must use unique names.

VOL=SER= - provide a dummy volume serial.

UNIT=(xxxx,,DEFER) - this allocates the proper tape drive type (xxxx) but does not try to mount the dummy volume serial.

DISP=(OLD,KEEP) - required

Only the first TAPEx DD provided is used by ABR during an Archive Restore.

EXAMPLE: // TAPE 1 DD DSN=FDR , VOL=SER=FDR ,
// UNIT=(TAPE , ,DEFER) , DISP=(OLD , KEEP)

DYNTAPE NOTE: If DYNTAPE is specified on the RESTORE statement, this DD statement is not used and can be omitted. ABR will dynamically allocate a TAPE# DD statement for the backup device. DYNTAPE should be used if the backup is on disk, in an automated tape library (ATL or silo) or if a mixture of tape device types must be read.

SIMULATION: If SIMREST is coded, this DD usually specifies DUMMY.

**ABRARCH DD
STATEMENT**

Specifies the remote queue data set for restores from the Archive Backups. This data set is optional. If specified, ABR will read the control statements contained within, if any, and append these control statements to the SYSIN data set. The SYSIN data set must contain at least a RESTORE TYPE=ARC statement. If the operation is not an restore from Archive Backup, this file will be ignored. After reading the control statements, ABR will reset the file to a null (empty) data set except on SIMREST.

NOTE: The TSO/ISPF panels or program FDRABRUT write to these remote queue data sets (see [Section 51.40](#)). DISP=SHR should be specified for ABRARCH, since ABR internally controls access to this data set.

**SYSIN DD
STATEMENT**

Specifies a data set containing the control statements for ABR. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

51.08 ARCHIVE RESTORE STATEMENT

RESTORE	TYPE=ARC	,ICFCAT=ORIGINAL STEP CAT ALIAS
R		
	,ALLOCATELIST=NO	,MAXCARDS=nnnn
SIMREST		
	,BLKF=nn	,NOCAT
		,RECAT
	,BYPASSACS	
		,OPERATOR
	,BYPASSSMS	
		,PRESTAGE
	,CATIFALLOC	
		,RESTORED=NO
	,COPY=n	
		,RLSE
	,DD=ALL	,%FREE=nn
	,DSNENQ=NONE TEST <u>USE</u> HAVE,	
		SELTERR=NO <u>YES</u>
	,DYNARC,	
		SMSGDG= <u>DEFERRED</u> ACTIVE ROLLEDOFF INPUT
	,DYNTAPE	
	,DYNTAPE2	,VRECAT

RESTORE STATEMENT The RESTORE TYPE=ARC statement activates a manual restore of individual data sets which have been archived. Only one RESTORE statement is allowed per execution of ABR, but any number of data sets can be restored in one ABR step.

SIMULATED RESTORE If SIMREST is coded, ABR will print the data set names which will be selected and the volumes necessary to do the restore. A restore operation is not done. It can be used to test RESTORE options or to pre-pull the tapes required.

51.08 CONTINUED . . .

**ARCHIVE
RESTORE
PROCEDURE**

The ARCHIVE RESTORE procedure restores data sets or groups of data sets that were archived by ABR. Since archived data sets usually do not exist on disk, ABR will allocate, catalog and restore the data set (although it will restore to a pre-allocated data set as well). See [Section 20.01](#) for details on how FDR allocates and catalogs data sets during restore.

ABR can also restore archived data sets automatically, when invoked from the ABR Catalog Locate exit (see the description of Auto-Recall in [Section 51.31](#)).

[Section 51.02](#) details how Archive Restore selects the appropriate Archive backup file on disk or tape for each data set to be restored. If multiple data sets are to be restored from a given backup, that file will be read only once. While reading a backup file, ABR can restore the selected data sets to one or more disk volumes concurrently. The target disk volume will be selected for each data set by the following rules:

- If the NVOL= operand was specified on the SELECT statement which selected this data set, that volume or volumes will be used. If NVOL= specified more than one volume serial, the first of those volumes will be selected initially; allocation may be attempted on up to 64 of those volumes in turn until it is successful. If the NVOL list includes more than one type of disk device, those with the same type as the input data set ("like" devices) will be tried first. Any volumes in the NVOL list which are not online will be ignored.
- If the output data set name is cataloged, then the volume to which it is cataloged will be chosen. The output data set name will be the original name or the new name if a NEWNAME=, NEWGROUP=, or NEWINDEX= operand was specified on the SELECT statement which selected the data set. If the data set is cataloged as being on multiple volume serials, then the volser will be selected from that list based on the volume sequence number in the F1 DSCB (field DS1VOLSQ) of the input data set.
- If none of the above apply, then the serial of the volume from which the data set was dumped, as recorded by ABR in the Archive Control File, will be used.
- If the data set was not preallocated on the selected volume, and the allocation fails on that volume for any reason, the ABR RESTORE ALLOCATION LIST, if enabled ([See Section 90](#)), will be checked to see if there is an ALLOCATE statement which applies to this data set. If so, the NVOL list from that statement will be used as described above for NVOL=. The RESTORE ALLOCATION LIST can be used to identify alternate volumes to which to restore data sets if their target volume is full or no longer exists. Note: you can force the RESTORE ALLOCATION LIST to be used for specific data sets by specifying NVOL=DUMMY or some other non-existent volume serial.

If SMS (System Managed Storage) is active on this system, and the data set does not already exist on the volume selected by the rules above, SMS is invoked to decide if the data set should be SMS-managed. If so, SMS will select an output volume. SMS rules are detailed in [Sections 20.01](#) and [70](#).

NOTE: if the records of certain archive data sets have been removed from the Archive Control File, ABR cannot automatically restore these data sets. If the backup files are still available, you can use FDRDSF to restore data sets from those backups.

**RESTORE ICF
VSAM FILES**

ABR will restore ICF VSAM files using the base cluster name. ABR will restore each individual component associated with this cluster name, and will allocate ICF VSAM files if they do not currently exist. ABR will update the appropriate fields within the VVR for each component. ICF VSAM files except the VVDS itself are movable; except for the VVDS and catalogs they can be restored to a new name or group. If NEWGROUP= or NEWINDEX= are specified the new group name will be applied to both the cluster name and all of its components. If NEWNAME= is specified for a cluster which is not allocated, ABR will let VSAM determine names for the components. Some information contained solely in the catalog, including protection (RACF or password) and expiration date will not be updated. However, path names for alternate indexes (AIXs) and aliases of user catalogs will be restored as long as both backup and restore are done with V5.3 level 30 or above. See [Section 80.13](#) "VSAM Special Considerations" for a more detailed explanation.

51.08 CONTINUED . . .

OPERANDS	TYPE=ARC	Required to restore data sets from ABR Archive Backups. ABR will attempt to restore all the data sets specified by the SELECT statements, DD=ALL option or remote queue data set.
	ALLOCATELIST=	<p>NO – specifies that the ABR RESTORE ALLOCATION LIST is not to be used for the selection of output volumes for the restore.</p> <p>Default is that the RESTORE ALLOCATION LIST will be used as described earlier in this section, if enabled in the FDR Global Option Table (see Section 90).</p>
	BLKF=	<p>Changes the blocksize (BLKSIZE) of certain fixed (RECFM=FB) and variable (RECFM=VB) data sets. BLKF= specifies the number of blocks per track (1 to 10) and is used to calculate a new blocksize based on the output track size (BLKF=1 results in BLKSIZE=32760). For RECFM=FB data sets the blocksize will be rounded down to a multiple of the logical record length (LRECL).</p> <p>DSORG=PS (sequential) data sets will be physically reblocked during the restore; the new blocksize must be larger than the original blocksize except when restoring from 3390 to 3380.</p> <p>DSORG=PO (partitioned) data sets will not be reblocked but the new blocksize will be used for new members; it must always be larger than the original blocksize.</p> <p>Other data set types, including UN (undefined), DA (direct access), PSE (enhanced format sequential), POE (PDSE), HFS and all types of VSAM, will not be reblocked.</p> <p>Default: data sets are not reblocked during restore; all original blocks will be restored without change, although they may be written to new locations. BLKF= is usually used when restoring to an unlike device type (e.g., 3380 to 3390) to improve utilization of the new track size but can also be used during like device restores.</p>
	BYPASSACS	<p>On a system with SMS (System Managed Storage) active, the SMS ACS (Automatic Class Selection) routines are not to be invoked for data sets which must be allocated. If a data set has a SMS storage class assigned (see STORCLAS= in Section 50.09) it will be SMS-managed, and SMS will be invoked to allocate the data set on an SMS-chosen volume, but SMS will not be allowed to override the storage class or management class assigned to the data set.</p> <p>Default: on an SMS system, the SMS ACS routines will be invoked for every data set which has to be allocated. The assigned storage and management classes will be passed to those routines, which can approve or override them. A data set will be passed to SMS for allocation if the storage class ACS routine assigns a storage class to the data set.</p>

51.08 CONTINUED . . .

BYPASSSMS

On a system with SMS (System Managed Storage) active, SMS data sets will be directly allocated on SMS-managed volumes, bypassing normal SMS storage group and volume selection. The selected output volume must be a SMS-managed disk volume, and the data sets being restored must have a SMS storage class assigned (see BYPASSACS above and STORCLAS= in [Section 50.09](#)). The data sets will be allocated and cataloged according to SMS standards.

Normal SMS facilities do not allow allocation of data sets on specific volume serials, but BYPASSSMS will do so, allowing data sets to be located for performance or other reasons. Note that if BYPASSACS is also specified, the assigned SMS classes will not be validity- or authority-checked.

Default: on an SMS system, for data sets which are SMS-managed and must be allocated, the SMS storage group ACS routine will be invoked to select a storage group and SMS will select a SMS-managed volume and allocate and catalog the data sets.

BYPASSACS and BYPASSSMS are primarily for use by storage administration personnel, since they bypass normal SMS allocation controls and rules. In order to use BYPASSACS or BYPASSSMS, the user of ABR must be authorized to the RACF profile

```
STGADMIN.ADR.RESTORE.BYPASSACS
```

in class FACILITY, or the equivalent in other security systems.

CATIFALLOC

non-VSAM output data sets will be cataloged even if they were preallocated (not allocated by ABR); the output data set will be cataloged if it is not already cataloged on another volume (unless the RECAT operand was specified).

Default: output data sets are cataloged only when the restore allocates them.

COPY=

Specifies the copy of the backup from which the restore is to be attempted; "n" can be 1 or 2. COPY=2 can be specified if a duplicate backup (TAPExx) was created at backup time.

Default is COPY=1 unless overridden in the FDR Global Option Table ([See Section 90](#)).

DD=ALL

ABR is to scan all the DISKxxxx DD statements for the DSNNAME coded. ABR will attempt to restore these data sets, unless they are temporary data sets, processing them as though they had been requested on a SELECT statement that specified DSN= with no other operands. DSNENQ=NONE is the default for these data sets.

51.08 CONTINUED . . .

DSNENQ=

Specifies whether all of the data sets being restored will be ENQed. See “Data Set Enqueue Option” in [Section 51.03](#) for more details.

If you are restoring over an existing data set and the ENQ fails, the restore will be bypassed with an error message. If the restore must allocate the output data set and the ENQ fails, no error message is issued and the restore is still done. A successful ENQ will prevent any other task from using the data set until the restore from the current backup data set is complete. An ENQ failure is considered an error but it will not prevent other data sets from being restored. The options for DSNENQ= are:

USE – The data sets will be enqueued for the duration of the restore from the current backup data set. This is the most frequently used option.

TEST – The data sets will only be tested to see if they are enqueued to another task at the time the restore starts. The data set will not be enqueued and other tasks may enqueue it while the restore is preceding.

HAVE – The data sets will be enqueued for the duration of the restore. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not enqueue the data set)

RETRY (try the enqueue again; may result in the FDRW27 message again)

NONE – No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: USE. Note that NONE or TEST may allow other jobs to attempt to read the data set being restored before ABR has restored all of the data tracks.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the restore is complete. However, use DSNENQ=NONE when another data set by the same name on another volume may be in use (e.g., restoring data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

**DYNTAPE
DYNTAPE2**

Specifies that ABR is to dynamically allocate the backup data sets using a DDNAME of TAPE#. This option should be used if the backup is on disk, in an automated tape library (ATL or silo) or a mix of backup device types will be read, such as 3480 and 3490E.

DYNTAPE2 will allocate 2 drives which will improve performance when restoring from multi-volume tape backups. DYNTAPE2 should not be used if backup files on disk might be involved.

Default: the first TAPEx DD statement found in the step JCL is used to mount all backups. The type of device assigned must be the correct type for reading all required backups.

51.08 CONTINUED . . .

ICFCAT=

Applies to ICF VSAM files only. Specifies the source of the catalog name to be used if an output ICF VSAM cluster must be allocated.

ORIGINAL – use the catalog in which the original dumped cluster was cataloged. When restoring a cluster to a new name, ICFCAT=ORIGINAL is treated like ICFCAT=ALIAS, described below. If you need to catalog the output cluster into the same catalog as the input cluster but that catalog is not the one aliased for the new name, you must specify ICFCAT=STPCAT and supply a STEPCAT DD statement pointing to that catalog.

STPCAT – use the STEPCAT as the target catalog. If a STEPCAT DD statement is not supplied, it will use the master catalog or the catalog which is aliased for this data set in the master catalog.

ALIAS – determine the catalog from the alias name in the master catalog. If no alias is found and the cluster is being restored to the same name, use the input cluster's original catalog. If no alias is found, and the cluster is being restored to a new name, it will use the STEPCAT (if present in the JCL) or the master catalog. Multi-level alias (MLA) is supported.

Default: ORIGINAL, except that if the cluster is being restored to a newname (NEWGROUP or NEWINDEX specified) the default is ALIAS. If the output cluster is SMS-managed, ALIAS is forced.

MAXCARDS=

Accept additional SELECT and EXCLUDE statements (over 100).

Default is 100 statements.

**NOCAT
RECAT**

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.

NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOP operand is also specified.

OPERATOR

Requests that an operator message (FDRW24) will be issued for each tape necessary to complete the restore. This option gives the operator the ability to pre-pull required tapes or bypass a tape which may not be available at this time.

51.08 CONTINUED . . .

PRESTAGE	<p>Output data sets which already exist on the target output volume will not be restored. This may be used to avoid restoring data sets which have already been restored. If the output data sets do not exist on the target volume, they will be allocated and restored.</p> <p>Default: pre-allocated data sets will have their contents overlaid.</p>
RESTORED=NO	<p>ABR will not restore data sets which are marked in the Archive Control File as having been previously restored.</p>
RLSE %FREE=	<p>RLSE – all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn – a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement).</p>
SELTERR=	<p>Specifies what will happen at step termination if one or more of the SELECT or EXCLUDE statements was never referenced (no data set was selected from the Archive Control File by the statement):</p> <p>NO – a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.</p> <p>YES – a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.</p> <p>Default: YES unless overridden in the FDR Global Option Table (See Section 90).</p>
SMSGDG=	<p>Specifies the status of SMS-managed GDG (Generation Data Group) data sets, if allocated by the restore.</p> <p>DEFERRED, ACTIVE, or ROLLEDOFF will set the GDG to that status.</p> <p>INPUT will set the GDG to the original SMS status of the GDG generation, as recorded on the backup tape. If the original GDG was non-SMS, it will be set ACTIVE if that generation is currently cataloged, otherwise DEFERRED.</p> <p>If a GDG is restored as DEFERRED, you may need to execute an IDCAMS "ALTER ROLLIN" to make the generation active.</p> <p>Default: DEFERRED, except that on automatic recall (other than via the remote queue) the default is INPUT.</p>

51.08 CONTINUED . . .

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when resotring a cluster when its catalog has been resotred, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume).

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for some reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

51.09 SELECT STATEMENT FOR ARCHIVE RESTORE

SELECT	DSN=filter	,NOTIFY=userid
S	DD=ddname	
	ALLDSN	,NVOL=(vvvvvv,vvvvvv,...)
EXCLUDE		
X	,ADATE=yyddd yyyyddd	,OLDBACKUP=nn
	,BLKF=nn	,PRESTAGE
	,COPY=n	,RESTORED=NO
	,DATA=ALL	,RLSE
		,%FREE=nn
	,DATACLAS=dataclass	
	,NULLDATACLAS	,STORCLAS=storageclass
		,NULLSTORCLAS
	,DSNENQ=NONE	
		,TRK=nnnnn
	,MGMTCLAS=managementclass	,CYL=nnnnn
	,NULLMGMTCLAS	
		,VOL=vvvvvv
	,NEWNAME=newdsname	
	,NEWGROUP=newgroup	,VRECAT
	,NEWINDEX=newindex	
	,NEWDD=ddname	
	,NOCAT	
	,RECAT	

SELECT STATEMENT This statement is used with RESTORE TYPE=ARC and selects the data sets to be restored from Archive Backups. The **SELECT** statement identifies an individual data set name or group of data sets to be restored. The **EXCLUDE** statement identifies data sets from within those selected by SELECT statements which are not to be processed. As described in [Section 51.02](#), ABR will select the backups of the selected data sets using records in the Archive Control File, locate the indicated backup on tape or disk, and restore it. EXCLUDE statements should only contain the operands DSN=, DD=, ALLDSN, or VOL=. A maximum of 100 of these control statements may be used in one execution unless overridden by MAXCARDS=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since ABR will only restore data sets which are selected, EXCLUDE statements are required only to exclude certain data sets from within a larger group on a SELECT statement.

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A . B . **
SELECT  DSN=A . **
```

Example 2: Select all data sets that were archived from volume TSO001 except those beginning with "ABC":

```
EXCLUDE DSN=ABC**
SELECT  ALLDSN , VOL=TSO001
```

51.09 CONTINUED . . .

**NEWNAME/
NEWGROUP/
NEWINDEX for
ICF VSAM** If you are restoring to a pre-allocated ICF VSAM cluster, where the cluster name is the same as the original, but the components may have different names from the original cluster, you must specify the cluster name as both DSN= and NEWNAME=; this causes ABR to locate the new component names. NEWNAME= cannot be specified for clusters with more than one alternate index.

If restoring an ICF VSAM cluster to a new name, if the new cluster must be allocated, you should specify NEWGROUP= or NEWINDEX=. ABR will modify both the cluster and component names.

[See Section 80.13 "VSAM Special Considerations".](#)

OPERANDS DSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the names of data sets on the backup tapes to be restored. For ICF VSAM clusters, only the cluster name is compared; you cannot select by component name.

```
EXAMPLES : DSN=USER1.JCL.CNTL
           DSN=**LIST
           DSN=PROD++.**.LIB*
```

You cannot use DSN= to select generation data group (GDG) generations by relative generation number. Use CATDSN= or DD= if relative GDG numbers are required.

Note: the CATDSN= operand will also be accepted on a SELECT for RESTORE, but it will be treated exactly like DSN=. It will not search system catalogs, nor will it accept GDG relative generation numbers.

DD=

Specifies that a data set name is to be taken from a DD statement. This operand must point to the DDNAME of a JCL statement. Using this option enables you to specify a non-standard data set name or a generation data set (GDG) relative generation.

```
EXAMPLE : SELECT      DD=DD1
                //DD1   DD   DSN=A.B.C(0),DISP=SHR
```

ALLDSN

By itself, ALLDSN will restore the most recent copy of **every data set** in the Archive Control File. You will usually use ALLDSN with ADATE=, to request that all data sets archived on the date specified are to be restored.

NOTE: DSN=, DD= and ALLDSN are mutually exclusive. One and only one of these operands must be specified on each SELECT and EXCLUDE statement.

ADATE=

Only data sets which were archived on the julian date specified (yyddd or yyyyddd) will be selected. Used to qualify a data set on the Archive Control File if the same data set name was archived multiple times. 2-digit year values (yy) less than 70 are assumed to be in the 21st century, e.g., ADATE=02123 means 2002.123.

If the data set was archived multiple times on the same day, OLDBACKUP= may also be specified to indicate which to restore.

51.09 CONTINUED . . .

BLKF=	<p>Selected PS and PO data sets are to be reblocked during the restore; see BLKF= in Section 51.08 for details.</p> <p>Default: data sets are not reblocked unless BLKF= was specified on the RESTORE statement. The restore will fail if the input data set has blocks larger than the track size of the output disk.</p>
COPY=	<p>Specifies the copy of the backup from which the restore is to be attempted. COPY=2 can only be specified if a duplicate backup (COPY 2) was created during archival or by FDRTCOPY.</p> <p>Default is COPY=1 unless COPY1 was not created or has expired, then COPY2 is the default if it exists. The default can be made COPY2 in the FDR Global Option Table (See Section 90) or by COPY=2 on the RESTORE statement.</p>
DATA=	<p>ALL – restore all allocated tracks in each data set selected by this SELECT statement.</p> <p>Default: restore only the used tracks of PS and PO data sets, unless DATA= was specified on the RESTORE statement.</p>
DATACLAS= NULLDATACLAS	<p>On a system with SMS active, specifies the SMS data class to be associated with the data set being restored, overriding the original data class of the data set (if any). The Data Class ACS routine will not be invoked.</p> <p>NULLDATACLAS changes the data class to null (not specified).</p> <p>Default: the original data class of the input data set (if any) will be associated with the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is set.</p>
DSNENQ=	<p>NONE – bypass the data set enqueue for data sets selected by this SELECT statement.</p> <p>Default: the enqueue option is determined by the DSNENQ= operand specified on the RESTORE statement.</p>
MGMTCLAS= NULLMGMTCLAS	<p>On a system with SMS active, specifies the SMS management class to be presented to the SMS Management Class ACS routine for the data set being restored, overriding the original management class of the data set (if any). The ACS routine may accept or override this class.</p> <p>NULLMGMTCLAS changes the management class to null (not specified).</p> <p>Default: the original management class of the input data set (if any) will be passed to the ACS routine for the output data set if it is allocated as SMS-managed. For a non-SMS input data set, a null class is passed.</p>
NEWNAME= NEWN=	<p>Restore the selected data set with a new name. NEWNAME should only be used with DSN= or DD=, and should not be used for ICF VSAM clusters unless they are preallocated. If the newname ends in a GDG relative generation number, e.g., NEWNAME=gdgname(-1), a LOCATE will be done to get the proper absolute generation number.</p>

51.09 CONTINUED . . .

**NEWGROUP=
NEWG=**

Restore the selected data sets using a new group name. The characters specified will replace the beginning of the input data set name. Care should be taken when periods are used that index levels are not incorrectly changed. ABR will check the new names for valid IBM standards. For ICF VSAM, the new group is applied to the cluster name and all component names.

EXAMPLE : SELECT DSN=ABC** , NEWG=XYZ

**NEWINDEX=
NEWI=**

Restore the selected data sets using a new name formed by adding or replacing one or more index levels in the original name; replacement index levels do not have to be the same length as the original indexes they replace. In the simplest case, each index level specified in NEWI is used in place of the corresponding index in the original name. Any remaining index levels at the end of the name are copied unchanged. This can easily be used to change the first indexes of the name.

For example, if the input data set is A.B.C.D,

NEWI=D results in D.B.C.D (first index replaced)

NEWI=DD.E results in DD.E.C.D (first 2 indexes replaced)

If a period is specified without any preceding characters, one original index level is copied from the input data set name to the output. This allows you to easily modify indexes in the middle of the name.

For example, if the input data set is A.B.C.D,

NEWI=..E results in A.B.E.D (third index replaced)

NEWI=FF...G results in FF.B.C.G (first and fourth indexes replaced)

If + is specified before a new index level, that new index is inserted into the output data set name at that point. If ++ precedes the new index, it will be added to the end of the name. If - is specified, the next input index level will be dropped from (not copied to) the output name.

For example, if the input data set is A.B.C.D,

NEWI=+F results in F.A.B.C.D (new first index added)

NEWI=..+F results in A.B.F.C.D (new third index added)

NEWI=++F results in A.B.C.D.F (new last index added)

NEWI=..- results in A.B.D (third index dropped)

NEWI=Q.-.+E results in Q.C.E.D (combination)

Note that, except for the ++ option, every period in the NEWI= mask corresponds to a period (one index level) in the original (input) data set name. The resulting new name will be checked to insure it meets IBM standards.

If the NEWI= value ends in a GDG relative generation number, e.g., NEWI=..NEWMAST(-2), that relative number will be added to the end of the newname, and a LOCATE done to get the proper absolute generation number.

NEWI= is a convenient way to rename every input data set, while using some index levels from the original name and replacing other indexes or adding new indexes. For ICF VSAM, the NEWINDEX is applied to the cluster name and all component names.

51.09 CONTINUED . . .

NEWDD=	<p>Specifies the name of a DD statement from which the new name of the output data set is obtained.</p> <p>NOTE: NEWN=, NEWG=, NEWI=, and NEWDD= are mutually exclusive. If none of them are specified, the data set is restored under its original name. NEWN= and NEWDD= should not be used on SELECT statements which select more than one data set.</p>
NOCAT RECAT	<p>NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.</p> <p>RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.</p> <p>Default: catalog output non-VSAM data sets only if they are not currently cataloged, unless overridden by NOCAT/RECAT on the RESTORE statement.</p> <p>NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.</p> <p>NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified on the RESTORE statement.</p>
NOTIFY=	<p>Specifies that ABR is to notify the TSO Userid specified at the completion of the restore. The userid is from 1 to 7 characters.</p>

51.09 CONTINUED . . .

NVOL=

Specifies the volume serial(s) of output disk volumes to which data sets selected by this statement are to be restored. You may specify:

- 1) A single specific volume serial, e.g.,
NVOL=ABC123
- 2) A list of specific volume serials, enclosed in parentheses, e.g.,
NVOL=(TSO001,TSO002,TSO003)
- 3) A group of volumes by placing an asterisk at the end of the volser prefix, e.g.,
NVOL=TSO*
- 4) A combination of specific and group, e.g.,
NVOL=(TSO*,PROD*,ABC001)
- 5) All online disk volumes may be selected by
NVOL=*

A list of online target volumes matching your specification is generated by scanning all disk UCBs in the system UCB chains; there is no guarantee of the order in which UCBs are found, so you cannot predict the order of the volume serials in the list. If you specify volume serials or groups which are not online, they are ignored and no error message will result.

However, if the first or only specification is a specific volume serial, it will be chosen as the first target volume, with other volumes placed after it in UCB chain order.

Also, if you are restoring a multi-volume data set (non-VSAM or SMS-managed VSAM), the volume sequence number of the piece of the data set being allocated will be used to select a specification from your list. For example, if NVOL=(A,B,C), the second piece of the data set will go to volume B. If that specification is a group, the first volume in the UCB chain matching that group will be tried. If the allocation is unsuccessful (such as insufficient free space), then other volumes in the NVOL list will be tried as described above.

The first target volume is checked to see if an output data set already exists there. If so, it restores over the existing allocation (unless PRESTAGE was specified). If not, it attempts to allocate the output data set on that volume. If the allocation fails, it will be retried on successive volumes in the list until it succeeds or until it fails on 64 volumes. If the list contains several disk device types, "like" volumes (same type as the data set being restored) will be tried first, then unlike devices.

For multi-volume data sets, a target volume is bypassed if a piece of the data set already exists there but is not the right piece, so that it will not attempt to restore the third volume of a data set on top of the first volume. When it finds a target volume in the list that does not contain a piece of the data set, it will be allocated.

Specifying multiple volsers or a volume group allows you to restore data sets in one pass even when no one volume has available space to contain them all; they will be spread across many of the target volumes.

Default: the output volume will be selected by rules defined in [Section 51.08](#). Note that when NVOL= is specified, and data sets are selected which are currently allocated and cataloged, ABR will restore them to the new volumes, and not to the volume on which they are cataloged.

On a system with SMS active, NVOL= may be ignored if the data set does not exist on the volume specified and the data set is SMS-managed (see STORCLAS= below).

Note: if an allocation is attempted on several volumes from your NVOL list but it fails on all of them, the message printed will usually show the allocation failure codes from the **first** volume only; failure codes from other volumes are not displayed and may be different.

51.09 CONTINUED . . .

OLDBACKUP=	<p>When a given data set has been archived more than once, specifies which versions to restore. 0 restores the most recently archived version, 1 the next most recent, etc., up to a maximum of 127. If ADATE= is also specified, OLDBACKUP= selects from multiple backups of the data set on that date only. (0 selects the last one archived on that date, etc.).</p> <p>Default is 0.</p>
PRESTAGE	<p>Specifies that selected data sets will not be restored if the output data set already exists on the first target output volume. This may be used to avoid restoring data sets which have already been restored.</p> <p>Default: restore pre-allocated data sets, overlaying the existing contents of those data sets, unless PRESTAGE was specified on the RESTORE statement.</p>
RESTORED=NO	<p>ABR will not restore a data sets if it is marked in the Archive Control File as having been previously restored.</p>
RLSE %FREE=	<p>RLSE – specifies that all of the unused space in the output PS (physical sequential) and PO (partitioned) data sets will be released.</p> <p>%FREE=nn – specifies a percentage (nn%) of the PS and PO data sets to be left free after the restore. However, the data sets will never be made larger than their original size. nn may range from zero (0) which will free all of the free space (same as RLSE) to 99 which attempt to leave the data sets with 99% free space.</p> <p>Space will be released only from data sets allocated by the restore; space is actually released by recalculating the required space during the allocation.</p> <p>Default: the output data sets are allocated the same size as the input data sets (unless overridden by TRK=/CYL= on the SELECT statement or by RLSE/%FREE= on the RESTORE statement).</p>
STORCLAS= NULLSTORCLAS	<p>On a system with SMS active, specifies the SMS storage class to be presented to the SMS Storage Class ACS routine for the data set being restored, overriding the original storage class of the data set (if any). The ACS routine may accept or override this class.</p> <p>NULLSTORCLAS changes the storage class to null (not specified).</p> <p>Default: the original storage class of the input data set (if any) will be passed to the ACS routine for the output data set. For a non-SMS input data set, a null class is passed.</p> <p>If the Storage Class ACS routine assigns a storage class to this data set or accepts the class passed, the data set will be allocated as SMS-managed, and the SMS Storage Group ACS routine may be invoked to determine the actual target volume. If the Storage Class ACS routine returns a null (blank) storage class name, the data set will be allocated as non-SMS and the ABR rules listed in Section 50.08 for volume selection must select a non-SMS target volume.</p>

51.09 CONTINUED . . .

TRK=
CYL=

If the data set selected by this SELECT statement must be allocated, CYL= or TRK= specifies the number of cylinders or tracks to be allocated to the data set. On PS or PO files when DATA=ALL was not specified, this value should be at least equal to the used portion of the data set. On all other types of files and when DATA=ALL is specified, this value should be equal to or greater than the original size of the file. For ICF VSAM clusters, modifies the size of the base data component only. If the space is too small for the data being restored, the restore will automatically extend the file for non-VSAM.

Default: use the original size of the data set.

VOL=

Specifies the disk volume serial number from which the data set name was archived. This operand is used to further delineate a data set on the Archive Control File if the data set name was archived multiple times.

(See [Section 51.02](#) for a more detailed explanation).

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume) unless VRECAT was specified on the RESTORE statement.

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for any reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

51.20 ARCHIVE BACKUP EXAMPLES

ARCHIVE ON LAST REFERENCE DATE FOR AUTO-RECALL Archive data sets which have not been opened in the last 30 days from two disk volumes, but only if they are not currently ENQed to any other task. Any data set on those volumes which was marked for archiving by the DISKUPDATE=YES option of the remote queue utility FDRABRUT will also be selected. The archived data sets will be recataloged for Auto-Recall with MIGRAT as the first or only volume serial. The backup tapes are to be kept for 180 days. A duplicate tape copy (TAPE11) is to be created. DYNARC will dynamically allocate the Archive Control File named in the FDR Global Option Table. At the end of the run the Archive Control File will be backed up as an ABR incremental Volume Backup.

```
//ARCHIVE      EXEC   PGM=FDRABR,REGION=3M
//SYSPRINT     DD      SYSOUT=*
//SYSPRIN1     DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//TAPE1        DD      DSN=ARC1,DISP=(,KEEP),UNIT=TAPE
//TAPE11       DD      DSN=ARC11,DISP=(,KEEP),UNIT=TAPE
//DISK1        DD      UNIT=DISK,VOL=SER=TSO123,DISP=OLD
//DISK2        DD      UNIT=DISK,VOL=SER=TSO124,DISP=OLD
//SYSIN        DD      *
DUMP           TYPE=ARC,ADAYS=30,RETPD=180,DYNARC,DSNENQ=USE,
               RECALL=YES,MIGRAT=YES
```

ARCHIVE ON ALL ONLINE VOLUMES Archive data sets on any online disk volume which have not been referenced in 60 days or which are not cataloged. Any archive requests which were added to the ABRARDQ remote queue data set by the DISKUPDATE=NO option of remote queue will also be selected. Any cataloged data sets which are selected by ADAYS=60 or the remote queue will be recataloged for Auto-Recall. The ABR Last Tape Option is used; if a previous ABR run used the same LASTAPE dsname, ABR will mount the last tape it used and add backup files to it. The backup files will be kept for 2 years (730 days). The backup of the Archive Control File will be done by DSF, creating file FDRABR.ARCBKUP at the end of the backup tape.

NOTE: uncataloged data sets are not eligible for Auto-Recall. However, if you have installed the ABR Data Set Not Found Exit, and JCL refers to the data set by UNIT= and VOL=SER=, that exit will recall the data set.

```
//ARCHIVE      EXEC   PGM=FDRABR,REGION=1200K
//SYSPRINT     DD      SYSOUT=*
//SYSPRIN1     DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//ARCHIVE      DD      DSN=FDRABR.ARCHIVE,DISP=SHR
//TAPE1        DD      DSN=FDRABR.LASTAPE.ARC1,
//              UNIT=TAPE,DISP=(MOD,KEEP)
//ABRARDQ      DD      DSN=FDRABR.ABRARDQ.DATA,DISP=SHR
//SYSIN        DD      *
DUMP           TYPE=ARC,ONLINE,ADAYS=60,IFNOTCAT,RETPD=730,
               ARCBKUP=DSF,DSNENQ=USE,RECALL=YES,MIGRAT=YES
```

51.20 CONTINUED . . .

**ARCHIVE
SPECIFIED
DATA SETS**

A group of data sets beginning with 'ABC' is to be unconditionally archived from disk volume 'TSO123'. However a subset of the data sets starting with 'ABC.DATA' is not to be archived. The ONLVOL option allows ABR to dynamically allocate the disk volume. All tracks of the data sets including the unused tracks are to be dumped. Two backup copies will be created using the Last Tape Option for each; COPY 1 will be retained for 120 days while COPY 2 is kept for 365 days (during restore, if COPY 1 has expired, COPY 2 is automatically used). The Archive Control File will **not** be backed up. The archived data sets will not be recataloged for Auto-Recall unless you have made RECALL=YES in the FDR Global Option Table (ISPF panel A.I.4.4).

```
//ARCHIVE      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSPRIN1     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//ARCHIVE      DD     DSN=FDRABR.ARCHIVE.MASTER,DISP=SHR
//TAPE1        DD     DSN=FDRABR.LASTAPE.ARC1,UNIT=TAPE,
//              DISP=(MOD,KEEP),LABEL=RETPD=120
//TAPE11       DD     DSN=FDRABR.LASTAPE.ARC11,UNIT=TAPE,
//              DISP=(MOD,KEEP),LABEL=RETPD=365
//SYSIN        DD     *
DUMP           TYPE=ARC,DATA=ALL,ONLVOL,ARCBACKUP=NONE,DSNENQ=USE
EXCLUDE        DSN=ABC.DATA**
SELECT         DSN=ABC**,VOL=TSO123
```

SIMULATION

SIMULATE an archive procedure on a group of disk volumes whose serials start with 'MAST', to verify that the proper data sets will be selected when the ARCHIVE is actually run. All data sets larger than 1000 tracks will be selected if they are not referenced in 10 days, data sets with 500 to 1000 tracks in 30 days, all others if they have not been referenced in 60 days.

```
//SIM          EXEC   PGM=FDRABR
//SYSPRINT     DD     SYSOUT=*
//SYSPRIN1     DD     SYSOUT=*
//TAPE1        DD     DUMMY
//SYSIN        DD     *
SIM            TYPE=ARC
SELECT         ALLDSN,SIZE=1001,ADAYS=10
SELECT         ALLDSN,SIZE=500,ADAYS=30
SELECT         ALLDSN,ADAYS=60
MOUNT         VOLG=MAST
```

**ARCHIVE
BASED ON
THRESHOLDS**

Archive data sets that have not been used in the last 30 days, but only from volumes where the percentage of tracks allocated exceeds a "low threshold" for the volume stored in the ABR Model DSCB or, for SMS-managed volumes, in the SMS storage group. The backup will be compressed by tape hardware compression (e.g., IDRC). The backup tape files will be retained for 365 days, by default.

```
//THRESHLD     EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT     DD     SYSOUT=*
//SYSPRIN1     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//TAPE1        DD     DSN=FDR,DISP=(,KEEP),UNIT=TAPE,DCB=TRTCH=COMP
//SYSIN        DD     *
DUMP           TYPE=ARC,ADAYS=30,THRESHOLD=LOW,DYNARC,DSNENQ=USE
MOUNT         VOLG=TST
```

51.21 AUTO RECALL AND POOLDISK EXAMPLES

Any of the examples in [Section 51.20](#) can be used to archive data sets to tape for Auto-Recall. Just be sure to specify RECALL=YES and MIGRAT=YES on the DUMP statement, or make both of those options the default in the FDR Global Option Table (ISPF panel A.I.4.4).

This section illustrates archiving to disk and tape, and the procedures necessary to maintain the output disk volumes. Typically, data sets are archived to disk, with a short retention period, so that auto-recall requests issued in the first few days after archive can be processed quickly without tape mounts. However, if you have an Automated Tape Library (ATL or silo) which can satisfy tape mounts quickly, you may want to archive only to tape.

ARCHIVE FOR AUTO RECALL

Archive data sets and recatalog them for Auto-Recall ([see Section 51.31](#)). DYNARC is used to allocate the standard Archive Control File, since it is the only one used for Auto-Recall. COPY 1 of the backup data set will be placed on a pool of disk volumes with a retention period of 30 days; it will be compressed by FDR. A duplicate copy will be created on 3490E scratch tapes with a retention of 1000 days; it will be compressed by the hardware compression feature (IDRC) of the tape drive.

```
//ARCHIVE      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSPRINT1    DD     SYSOUT=*
//TAPE1        DD     DSN=FDRABR.POOLDISK.ARCHIVE1,
//              DISP=OLD,UNIT=DISK,LABEL=RETPD=30,
//              VOL=SER=(ARC100,ARC101,ARC102,ARC103)
//TAPE11       DD     DSN=FDR11,DISP=(,KEEP),UNIT=3490,
//              LABEL=RETPD=1000,DCB=TRTCH=COMP
//SYSIN        DD     *
DUMP           TYPE=ARC,ONLINE,DYNARC,RECALL=YES,MIGRAT=YES,
              ADAYS=60,COMPRESS=COPY1,DSNENQ=USE,RTC=YES
```

In this example, the COPY2 files created on TAPE11 might be sent off-site for disaster recovery. At the disaster site, to cause ABR to automatically use the off-site tapes for restores instead of the POOLDISK files, change option ARCCOPY to "2" in the FDR Global Option Table ([see Section 90](#)).

For on-site restores, ABR will automatically use COPY1 on POOLDISK until its expiration date (30 days) is reached, then it will automatically try to use COPY2 on tape. If the COPY2 tapes are off-site, or if you want to keep 2 copies for on-site restore after the migrated copy on disk expires use the following FDRTSEL example.

MOVE DISK ARCHIVES TO TAPE

Use FDRTSEL to select Archive Backup files on POOLDISK (COPY 1) that will expire within the next 3 days, moving them to tape. The expiration of the COPY 1 on tape is 700 days from now. The Archive Control File will be updated to point to the new COPY 1, and the disk backup files will be scratched. [See Section 60](#) for more details on FDRTSEL.

```
//FDRTSEL      EXEC   PGM=FDRTSEL,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//TAPEOUT      DD     DSN=TSEL1,DISP=(NEW,KEEP),UNIT=TAPE,
//              LABEL=RETPD=700
//SYSIN        DD     *
MOVE          CAT=RECAT,DYNARC
SELECT        ARCHIVE,COPY=1,EXPIRE=YES,XDAYS=3
```

CATALOG POOLDISK DATA SET

This is an example of using IDCAMS to create a catalog entry for an ABR POOLDISK. This allows the POOLDISK to be referred to by name without having to specify the disk volumes in the ABR JCL. A DELETE is included to delete any previous definition. This is only a catalog entry, no actual data set is created.

```
//POOLCAT      EXEC   PGM=IDCAMS
//SYSPRINT     DD     SYSOUT=*
//SYSIN        DD     *
DELETE FDRABR.POOLDISK.ARC1 NOSCRATCH
SET MAXCC=0
DEFINE NONVSAM(NAME(FDRABR.POOLDISK.ARC1) -
  DEVICETYPE(3390) -
  VOLUMES(ARC001 ARC002 ARC003))
```

51.21 CONTINUED . . .

CLEAN UP DISK ABR BACKUPS Clean up expired ABR backup files from two POOLs of disks. ABR will search the volumes specified for data sets which meet ABR's naming conventions for backup data sets. If they have expired, ABR will scratch these data sets without taking a backup. The second POOL is referenced as a cataloged data set.

This job should be run occasionally if you are doing ABR Archive Backups or Volume Backups to disk, to clean up expired files and maintain as much free space as possible on those disks. However, since FDRTSEL can copy and scratch the disk backups (see example above) and FDRARCH REORG will scratch expired disk backups, this job will usually select only data sets which were not scratched for some reason.

```
//SCRATCH      EXEC   PGM=FDRABR
//SYSPRINT      DD     SYSOUT=*
//SYSPRINT1     DD     SYSOUT=*
//TAPE1         DD     DUMMY
//DISK1         DD     DSN=FDRABR.POOLDISK.ARCHIVE1,DISP=OLD,
//              VOL=SER=(ARC100,ARC101,ARC102,ARC103),UNIT=DISK
//DISK2         DD     DSN=FDRABR.POOLDISK.ARCHIVE2,DISP=OLD
//SYSIN         DD     *
                DUMP    TYPE=SCR,SCRATCH
                SELECT  ABRBKUP
```

51.22 ARCHIVE RESTORE EXAMPLES

RESTORE A DATA SET GROUP Restore a group of data sets with the high-level index of 'USERxx'. The most recent copy of each such data set recorded in the standard Archive Control File will be restored. DYNTAPE is used to allocate the backup tape or disk. The data sets will be restored to their original volumes; if any cannot be allocated on that volume, the ABR RESTORE ALLOCATION LIST, if enabled, will be used to choose alternate volumes.

```
//RESTORE      EXEC   PGM=FDRA BR, REGION=1024K
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
      RESTORE    TYPE=ARC, DYNARC, DYNTAPE
      SELECT    DSN=USER%.**
```

RESTORE SELECTED DATA SETS Specified data sets are to be restored from Archive Backup. The operator will be notified of the tape volumes required before a restore is attempted on each tape. All data sets with a group name of 'MARK1.' will be restored. Also data set 'OLDFILE', which was archived on 97.350 from disk volume SYS002, will be restored to data set 'NEWFILE' on disk volume SCR011. The third oldest archive version of data set 'MYFILE' will be restored to volume SYS001.

Although DYNTAPE is the preferred way of allocating the backup files to be read, this example uses the alternate method of providing a TAPEX DD statement. It must allocate a device capable of mounting and reading all of the required backups (3480, 3490E, 3590, Redwood, 9840, 3380, 3390). In this example, all of the archive backups are on 3480 cartridge, so a TAPE1 DD is provided pointing to a 3480 drive.

```
//RESTORE      EXEC   PGM=FDRA BR, REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRA BR. ARCHIVE, DISP=SHR
//TAPE1         DD     DSN=FDRA BR, VOL=SER=FDRA VOL,
//              UNIT=(3480,,DEFER), DISP=(OLD,KEEP)
//SYSIN         DD     *
      RESTORE    TYPE=ARC, OPERATOR
      SELECT    DSN=MARK1.**
      SELECT    DSN=OLDFILE, VOL=SYS002, ADATE=97350,
              NEWN=NEWFILE, NVOL=SCR011
      SELECT    DSN=MYFILE, NVOL=SYS001, OLDBACKUP=2
```

REMOTE QUEUE RESTORE A job similar to the following should be run several times a day if the ABR Archive restore remote queue data set is used. It will process any queued requests and empty the queue so that more requests can be added. The frequency of this job must be decided by the installation, based on how long users are willing to wait for their restore request to be completed (1 hour, 4 hours, etc.).

```
//REMTAE       EXEC   PGM=FDRA BR, REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRA BR. ARCHIVE, DISP=SHR
//SYSIN         DD     *
      RESTORE    TYPE=ARC, DYNTAPE
//ABRARCH      DD     DSN=FDRA BR. ABRARCH. USER, DISP=SHR
```

51.22 CONTINUED . . .

RESTORE BY Restore all of the data sets archived on the date 98.123 from a 3380 volume 'PRODUC' to an unlike
ARCHIVE device (3390). All but 10% of the unused space within PS and PO data sets will be released. All
DATE data sets will be recataloged to the new volume.

```
//UNLIKE      EXEC  PGM=FDRABR,REGION=2M
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSUDUMP    DD    SYSOUT=*
//SYSIN       DD    *
      RESTORE  TYPE=ARC,%FREE=10,DYNTAPE,DYNARC,RECAT,VRECAT
      SELECT   ALLDSN,NVOL=V33901,ADATE=98123,VOL=PRODUC
```

YEAR2000 Note: In the ADATE= operand, year numbers less than 70 will be interpreted as dates in the 21st century. For example, ADATE=01222 is 2001.222.

51.23 SUPERSCRATCH EXAMPLES

NOTE: any volume to be processed for Superscratch must be enabled for Superscratch in the ABR Model DSCB on that volume (ISPF panel A.I.8).

SIMULATE SUPER-SCRATCH Because Superscratch can delete many needed data sets if it is misused, **you should always run a simulation** of any new Superscratch jobstream and verify that it selected only the desired data sets. This is a simulation of the next example.

```
//SCRATCH      EXEC   PGM=FDRA BR , REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSPRIN1      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DUMMY
//SYSIN         DD     *
SIM             TYPE=SCR , IFNOTCAT , SCRATCH=CAT
MOUNT           VOLG=TSO
MOUNT           VOLG=PROD
```

SCRATCH UNCATALOGED DATA SETS All data sets in the installation must be cataloged, so all uncataloged data sets on TSO and PROD volumes are to be scratched if they are not properly cataloged. The IBM SCRATCH SVC will be used to scratch the data sets; any scratched data set which has a current ABR Volume Backup will be recorded in the ABR scratch catalog, so it can be recovered from its backup if required.

```
//SCRATCH      EXEC   PGM=FDRA BR , REGION=1M
//SYSPRINT      DD     SYSOUT=*
//SYSPRIN1      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DUMMY
//SYSIN         DD     *
DUMP            TYPE=SCR , IFNOTCAT , SCRATCH=CAT
MOUNT           VOLG=TSO
MOUNT           VOLG=PROD
```

NOTE: if you want to select only certain uncataloged data sets, remove IFNOTCAT from the DUMP statement and insert a SELECT statement with at least DSN= and IFNOTCAT, e.g.,

```
DUMP            TYPE=SCR , SCRATCH=CAT
SELECT          DSN=XYZ** , IFNOTCAT
MOUNT           VOLG=TSO
```

or

```
DUMP            TYPE=SCR , SCRATCH=CAT , ONLVOL
SELECT          DSN=XYZ** , IFNOTCAT , VOLG=TSO
```

A common mistake is to use statements such as

```
DUMP            TYPE=SCR , IFNOTCAT , SCRATCH=CAT , ONLVOL
SELECT          ALLDSN , VOLG=TSO
```

Since the SELECT statement overrides IFNOTCAT on the DUMP statement, **all data sets on the TSO volumes will be archived** (see details in [Section 51.02](#)).

CLEAN OFF WORK VOLUMES All WORK volumes are to be cleaned off nightly. DSNENQ=USE ensures that any data sets still in use will not be selected. TEMP will select temporary data sets; ALLDSN selects all other data sets.

```
//SCRATCH      EXEC   PGM=FDRA BR
//SYSPRINT      DD     SYSOUT=*
//SYSPRIN1      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPE1         DD     DUMMY
//SYSIN         DD     *
DUMP            TYPE=SCR , ONLVOL , DSNENQ=USE
SELECT          TEMP , VOLG=WORK , CRDAYS=2
SELECT          ALLDSN , VOLG=WORK
```

51.24 SMS ARCHIVE EXAMPLES

ABR Archive and Superscratch on SMS-managed volumes can be done in the same ABR step as non-SMS volumes. Selection of data sets from SMS volumes can be done by normal ABR control statements (if SMSMANAGE=NO is specified or defaulted), or by SMS management class criteria (SMSMANAGE=YES must be specified). See Section 70 for more details. The examples below process only SMS volumes so that SMS processing can be clarified.

**SCRATCH AND
ARCHIVE
BASED ON
MANAGE-
MENT CLASS**

Since some of the characteristics of the SMS management class relate to scratching data sets, and some to archiving, a complete implementation of the SMS management class requires that a ABR Superscratch step be run right before a ABR Archive step. This is explained in detail in Section 70. In this example, only the management class criteria will be used to select data sets for scratch and Archive. **Note that the ABR RESTORE PROTECT LIST will not prevent selection of SMS data sets.**

```
//SCRATCH      EXEC  PGM=FDRABR,REGION=1200K
//SYSPRINT      DD    SYSOUT=*
//SYSPRIN1      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//TAPE1         DD    DUMMY
//SYSIN         DD    *
      DUMP      TYPE=SCR,SMSMANAGE=YES,DSNENQ=USE,EXPIRED,SELTERR=NO
      MOUNT      STORGRP=TSO1
      MOUNT      STORGRP=TSO2
//ARCHIVE      EXEC  PGM=FDRABR,REGION=2148K,COND=(0,NE,SCRATCH)
//SYSPRINT      DD    SYSOUT=*
//SYSPRIN1      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//TAPE1         DD    DSN=FDRABR.POOLDISK.SMS1,UNIT=DISK,
//              VOL=SER=(ARC001,ARC002,ARC003),DISP=OLD,LABEL=RETPD=30
//TAPE11        DD    DSN=SMS2,UNIT=TAPE,DISP=(,KEEP),LABEL=RETPD=730
//SYSIN         DD    *
      DUMP      TYPE=ARC,SMSMANAGE=YES,RECALL=YES,MIGRAT=YES,DSNENQ=USE,
      COMPRESS=COPY1,DYNARC
      MOUNT      STORGRP=TSO1
      MOUNT      STORGRP=TSO2
```

**ARCHIVE
BASED ON SMS
THRESHOLDS**

Archive data sets from SMS-managed volumes based on their SMS management class criteria, but only from SMS volumes whose percentage of tracks allocated exceeds the high threshold associated with the volume's SMS storage group.

```
//ARCHIVE      EXEC  PGM=FDRABR,REGION=2248K
//SYSPRINT      DD    SYSOUT=*
//SYSPRIN1      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//TAPE1         DD    DSN=SMS1,UNIT=TAPE,DISP=(,KEEP),LABEL=RETPD=60
//TAPE11        DD    DSN=SMS2,UNIT=TAPE,DISP=(,KEEP),LABEL=RETPD=365
//SYSIN         DD    *
      DUMP      TYPE=ARC,SMSMANAGE=YES,RECALL=YES,MIGRAT=YES,DSNENQ=USE,
      THRESHOLD=HIGH,DYNARC
      MOUNT      STORGRP=PROD1
```

NOTE: if the MAXBTRKS= operand is used to limit the number of data tracks recorded in a single backup file (default: 4095), the threshold is reevaluated at the beginning of each pass on the volume, so it will bypass the volume if the previous pass reduced the volume allocation below the threshold. If you want this to occur, you might specify a smaller value, e.g., MAXBTRKS=500.

51.30 ARCHIVE BACKUP ISPF SUPPORT

The ABR ISPF dialogs support data set archives and restores for Archive Backups. Users can use them to submit their own ABR jobs to restore specific data sets from the Archive Backups, or they can add the requests to a Remote Queue data set (described in [Section 51.03](#)).

ABR ISPF MAIN MENU

Innovation recommends that the ISPF dialogs be installed so that a character of "A" on your regular ISPF main menu will call the ABR main menu. Details of this installation are found in [Section 90](#). When this manual refers to a panel name such as A.S, it assumes that this convention was used. If your installation used a different letter on your ISPF main menu, or uses some other technique to get to the ABR main menu, you will have to change our examples appropriately.

If you enter simply "A" (or whatever) from your ISPF main menu, the ABR main menu is displayed:

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- FDR PRIMARY OPTIONS MENU -----
OPTION  ===>

    1  REPORTS      - ABR REPORTING FUNCTIONS
    2  RESTORE      - ABR DATA SET RESTORE
    3  ARCHIVE      - ABR DATA SET ARCHIVE OR SUPERSCRATCH
    4  BACKUP       - ABR DATA SET BACKUP
    5  REMOTE Q     - ABR REMOTE QUEUE UTILITY FUNCTIONS

    C  COMPAKTOR    - COMPAKTOR MAP AND SIMULATION REPORTS
    R  RELEASE      - COMPAKTOR RELEASE

    I  INSTALL      - INSTALLATION AND MAINTENANCE OF FDR AND OPTIONAL PRODUCTS
    J  JCL PARMS    - SPECIFY FDR JCL AND SYSOUT DEFAULTS FOR SUBMITTED JOBS
    K  FORMAT       - MODIFY FORMAT OF GENERATED REPORTS

    M  MESSAGES     - FDR MESSAGES ADN CODES QUERY FACILITY
    Q  QUERY        - FDR/ABR STATISTICS QUERY

    S  SRS          - SEARCH, REPORT, SERVICES DIALOG

    T  FDRTSEL      - BACKUP FILE MANAGEMENT UTILITY

```

Note that option 2 (RESTORE) can be used to restore the archived data sets from Archive Backups, while option 3 (ARCHIVE) can request that certain data sets be archived. Depending on the customization of the dialogs done during install ([Section 90](#)), the user may have the option to directly submit an ABR job or to add the request to a remote queue. Although the RESTORE panels supports restore of older backups, the SRS dialog, described next, is a far superior way of restoring data sets.

SRS ISPF DIALOG

SRS stands for "Search, Report, and Services". It is described in detail in [Section 54](#). It allows users to search for information on data sets from a variety of sources, display the information in a easy-to-use format, and optionally to execute various TSO and ABR functions against the data sets displayed.

For Archive Backups, it can display backup information for the data sets selected. If data sets have been archived more than once, it can display all of the recorded backups, including the date of the archive. Simply by typing the RESTORE command on the proper line, users can request that a data set be restored.

Users can also add a request to the backup Remote Queue, so that the displayed data set will be included in the next incremental backup of the volume it resides on.

51.30 CONTINUED . . .

From your ISPF main menu, if you enter "A.S.1", you will receive this SRS panel:

----- DATASET SELECTION: ABRBKUP ----- LINE 1/16 COL 4:6/10									
COMMAND ==>					SCROLL ==> HALF				
ENTER SELECTION CRITERIA									
	Read	Save	Submit	Find	Locate	Extract	Options	Help	
FIELD	SELECTION VALUE							REPORT	SORT
-----	-----							-----	----
DSNAME	==>	bab.multi.**_						----	----
VOL	==>	-----						2_	----
SOURCE	==>	ARCHIVE (Catalog Volume Archive Appl Scratch Extract)						----	----
ARCDSN	==>	-----							
DEVTYPE	==>	-----						3_	----
DSORG	==>	-----						4_	----
BKINFO	==>	-----						6_	----
SIZE	==>	-----						5	

In this example, the user has filled in a data set name mask, has specified several attributes of the selected data sets to display (including ABR backup information), and had requested that all recorded backups of each data set be displayed (OLDBACKUP=ALL). Information from the Archive Control File has been requested (ARCDSN will default to the common Archive Control File named in the FDR Global Option Table).

Note that all of the SRS report and selection fields shown may not display on the screen at once, you may need to scroll up and down to find the SRS fields you need. [Section 54](#) includes information and examples on customizing SRS to display just the information required for a particular function, such as restore from Archive Backup.

In response to this request, SRS will search the Archive Control File for the data sets requested and display their backup information in a format similar to:

----- DATASET LIST: DSLIST ----- LINE 1/42 COL 3:7/13									
COMMAND ===>					SCROLL ===> HALF				
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help
COMMAND	ENTRY NAME				VOLSER	DEVTYPE	DSO	ALLOC	BKDATE
-----	-----				-----	-----	-----	-----	-----
	BAB.MULTI.VOLTEST				SMS812	3380	**	4	1997.234
	BAB.MULTI.VOLTEST				SMS811	3380	PS	16	1997.234
restore	BAB.MULTI.VOLTEST				SMS812	3380	**	4	1997.210
restore	BAB.MULTI.VOLTEST				SMS811	3380	PS	16	1997.210

For data set BAB.MULTI.VOLTEST, it was archived from volumes SMS812 and SMS811 on 97.210, and again on 97.234. If you want to restore the version from 97.210, enter the RESTORE command next to both pieces of the data set.

51.30 CONTINUED . . .

The RESTORE command will take you to another panel where you can set options for the restore:

```

----- FDRSRS - Archive Restore ----- Row 1 to 2 of 2
COMMAND ==> submit                                SCROLL ==> HALF

  Edit JCL    Submit JCL    FG - execute in the foreground    RQ - add to remote q

Operands for RESTORE TYPE=ARC statement (section 51.08):
==> RESTORE TYPE=ARC,DT,ARCDSN=FDRABR.ARCHIVE

DSNAME / Filter ==> 'BAB.MULTI.VOLTEST'
Volume Serial   ==> SMS812                                Archive date ==> 97210
New DSNNAME     ==>
or NEWINDEX     ==>
New Volser(s)   ==>                                Copy ==> 1
Operands for SELECT DSN= statement (section 51.09):
==> NOTIFY=BAB

DSNAME / Filter ==> 'BAB.MULTI.VOLTEST'
Volume Serial   ==> SMS811                                Archive date ==> 97210
New DSNNAME     ==>
or NEWINDEX     ==>
New Volser(s)   ==>                                Copy ==> 1
Operands for SELECT DSN= statement (section 51.09):
==> NOTIFY=BAB

```

The proper volume and archive date have already been filled in. If required, you can enter a new name or NEWINDEX for renaming the restored data set; you can also specify a new target volser if it is not to be restored to the original volume.

The user has the option to submit a batch ABR jobstream to perform the restore (EDIT or SUBMIT commands), to attach ABR under TSO and execute the restore in the foreground (FG command), or to add the requests to the ABR Restore Remote Queue (RQ command, only valid if you have implemented that remote queue as described in [Section 90](#)). Note that foreground (FG) restores from tape will work only if the TSO user is authorized to mount tapes. Your installation may have modified this panel to limit the choices available, or it may execute a particular command (e.g., RQ) automatically without displaying this panel.

Because of the NOTIFY= operand, your TSO session will receive a message when the restore of each data set is complete.

51.31 ABR AUTO-RECALL - INTRODUCTION

As described in [Section 51.01](#), ABR Archive Backups can select data sets from disk and move them to a less expensive medium such as tape or compressed disk. The usual criteria for selection of data sets to be archived is the Last Reference Date of each data set, selecting data sets which have not been used recently.

But no matter what criteria you use for selection, some of those archived data sets will eventually be required by jobs or TSO users. This might be due to jobs which run at infrequent intervals, test data sets for applications which have not been tested recently, TSO libraries which are infrequently used, or TSO users who have not logged on recently.

It is possible to manually restore archived data sets as described in [Section 51.07](#), if you can predict which data sets will be required. But it will be much more convenient if the required data sets are automatically restored when they are needed, transparent to the requesting task. This is the function of ABR Automatic Recall (Auto-recall). When properly implemented, Auto-Recall will dynamically recall archived data sets referenced by:

- DD statements in JCL. This includes batch jobs, started task JCL, and TSO startup procs. Archived data sets are recalled at the beginning of each step, before the job step program gets control; all archived data sets in the step are recalled at one time. The step is delayed until the recall is complete, so when the program executes all required data sets have been placed back on disk. Other than ABR recall messages in the job log, the job is unaware that they were ever archived.
- dynamic allocation by a batch job, started task or TSO user. For a batch job, the data set being allocated is recalled immediately; the allocation will wait until the recall is complete. For a TSO user, the user is notified that an archived data set is requested, and given the choice of exiting, waiting for the recall to complete, or recalling the data set without waiting, allowing other work to be done while the recall is proceeding. You can customize ABR to limit or eliminate these choices.

Auto-Recall is performed by two Operating System exits that are supplied with the ABR system: the Catalog Locate Exit and the Data Set Not Found Exit. [See Section 90](#) and ISPF panel A.I.4.11 for installation and customization of these exits.

**CATALOG
LOCATE EXIT**

The ABR Catalog Locate Exit intercepts catalog requests (such as LOCATE SVCs) issued by system components (such as JCL allocation and dynamic allocation), TSO components and programs, third-party software and even user-written programs. If the requested data set is cataloged, and the catalog entry indicates that the data set has been archived with the RECALL option, then the Catalog Locate Exit arranges for FDRABR to restore the data set. The Catalog Locate Exit may invoke FDRABR directly within the address space requesting the data set, or may issue a START command to invoke FDRABR in a separate address space, or may place the restore request into the remote queue data set to be performed along with other ARCHIVE restores. Details appear in the following Sections.

**DATA SET NOT
FOUND EXIT**

The Data Set Not Found Exit receives control when a program issues an OPEN for a disk data set and OPEN cannot find that data set in the VTOC of the volume specified. Since the Catalog Locate Exit will recall data sets when a catalog LOCATE is done for them, they will normally be recalled long before OPEN is done. So, this exit is usually invoked only when a volume serial is specified in batch JCL, since the catalog search is not done. For example:

```
//DD1 DD DSN=TEST.FILE,UNIT=SYSDA,VOL=SER=TEST01,DISP=OLD
```

If your installation's standard is that all data sets must be cataloged, the Data Set Not Found exit will be rarely used and you may choose not to implement it at all.

The Data Set Not Found Exit **always** invokes FDRABR in the address space of the job issuing OPEN to see if the data set has been archived (recorded in the Archive Control File as having been archived from the volume involved); the data set need not be cataloged for auto-recall. If it was archived, it is recalled to that volume and OPENed. If ABR does not find it, OPEN will issue a normal

51.31 CONTINUED . . .

S213-04 ABEND as if ABR was never involved. If FDRABR internally ABENDs during the recall, OPEN will issue an S213-2C ABEND.

NOTE: When the Data Set Not Found Exit is installed, any S213-04 ABEND will be preceded by the message

FDR316 RECALL DID NOT FIND REQUESTED DSN/DSG=dsname

Many users think that this means that an ABR error has occurred, but in fact this simply means that the data set being OPENed has not been archived. The S213-04 ABEND means that the data set was truly not found on disk. This could be due to a misspelling of the data set name or the wrong volume serial. This will also occur if a DD statement points to a VSAM cluster, but a non-VSAM OPEN is being done by the program.

ARCHIVE CONTROL FILE Auto-Recall requires that all recallable data sets be recorded in a common Archive Control File, whose name is specified in the FDR Global Option Table. This Archive Control File will be dynamically allocated and searched for all Auto-Recall requests.

RECALL INDICATOR When a data set is archived with RECALL=NO, its catalog entry will be deleted (for multi-volume data sets, this is done when all pieces of the data set have been archived).

But when a data set is archived for auto-recall (the RECALL=YES operand specified or defaulted), archived data sets remain cataloged, with a special indicator in the catalog entry that tells the Catalog Locate Exit (see below) that the data set is recallable. This indicator also includes an index into the Archive Control File so that recalls can begin reading the control file near the record for the data set, rather than scanning the entire file.

For single-volume data sets, this indicator is stored in a catalog field called DSCBTTR. For data sets on disk, this field often points to the Format 1 DSCB of the data set, to speed up OPEN processing. Since there is no DSCB for archived data sets, ABR uses it to store the auto-recall information. DSCBTTR is not displayed by LISTCAT and similar catalog list programs, but FDREPORT ([Section 54](#)) can do so.

For multi-volume data sets, since each piece could be archived separately, the auto-recall indicator is stored in the File Sequence field of each volume in the catalog entry. The Catalog Locate exit will recall only the pieces which are marked for auto-recall. LISTCAT will display the File Sequence field in decimal, so you may see some large values; this is normal.

For VSAM clusters, the VSAM catalog entry is deleted when the cluster is archived, and a non-VSAM catalog entry is built with the name of the cluster, the volumes it was archived from, and the auto-recall indicator. For DB2 files, the data component name is also cataloged for Auto-Recall, since DB2 may refer to the component name. For multi-volume VSAM, the cluster is not deleted until the cluster has been archived from every disk volume on which it resides, so the non-VSAM entry is built when the cluster is finally deleted, with all the volumes and auto-recall indicators set. **To properly archive multi-volume VSAM, all volumes containing pieces of the cluster must be processed in the same ABR run.**

RECALL=YES can be made the default in the FDR Global Option Table (ISPF panel A.I.4.4). This is recommended if you plan to archive all or most data sets for Auto-Recall. RECALL=NO can override it for specific jobs.

MIGRAT OPTION ABR has an option to change the first volume serial in the catalog entry for an archived data set to MIGRAT; this is the indicator used by IBM to indicate that a data set has been migrated (archived) by DFHSM or DFSMSHsm. ABR supports the MIGRAT volser so that users can more easily identify archived data sets from a LISTCAT or other data set display. Because of checks in IBM code for the MIGRAT volume serial, it also avoids some errors which can occur since IBM does not recognize our auto-recall indicator.

When RECALL=YES is used, MIGRAT=YES can also be specified (or defaulted in the FDR Global Option Table, ISPF panel A.I.4.4) to change the first or only volume serial to MIGRAT.

Warning: MIGRAT=YES is required when archiving SMS-managed data sets and is recommended for ALL data sets.

51.32 ABR AUTO-RECALL - IMPLEMENTATION

The following is a list of the steps an installation should take to implement Auto-Recall:

- STEP 1** Install the Catalog Locate Exit and, optionally, the Data Set Not Found Exit, as described in [Section 90](#). When new releases of ABR are installed, the exits distributed with that release should be installed at the same time; [Section 90](#) also has details on testing new versions of the exits. Please read [Section 51.34](#) to select customization options for the exits.
- STEP 2** Create the common Archive Control File to be used for all auto-recallable data sets, as described in [Section 90](#).
- STEP 3** When a recall is done, ABR updates the Archive Control File to flag the data set has having been restored. If your installation has a security system, you may need to give authorization to all users to UPDATE the Archive Control File; if possible, allow UPDATE access only from a program whose name starts with the 'FDR' that was loaded from the FDR program library. You may also need to grant users authority to READ the Archive Backup files on tape or disk. For recalls by the Catalog Locate exit, these security rules are not necessary if the LXBYPSEC option is enabled (see [Section 51.34](#)).
- STEP 4** You must schedule Archive Backup jobs to archive data sets for auto-recall. Many installations do this daily, but you might choose to do it weekly, monthly, or even multiple times per day. Your ABR Archive Backup job should specify these operands on the DUMP TYPE=ARC statement:
- **RECALL=YES** - to recatalog the archived data sets for Auto-Recall.
 - **MIGRAT=YES** - to change the first volser in the Auto-Recall catalog entry to MIGRAT (required for SMS-managed data sets).

Both of these options are described in [Section 51.31](#) and [51.05](#) and each can be made the default in the FDR Global Option Table (ISPF panels A.I.4, see [Section 90](#)). You can override the defaults if you have Archive Backups that do not require auto-recall (e.g., RECALL=NO) so Innovation recommends that you set these in the option table if you use auto-recall for most data sets.

- STEP 5** In many cases, the Archive Backup JCL will direct the first copy (COPY 1) of the Archive Backups to disk, and a second copy (COPY 2) to tape. The disk backup will usually have a short retention (e.g., 15-30 days) while the tape backup will be retained for as long as your installation requires. This allows recalls that occur shortly after the data set is archived to be quickly satisfied from the disk copy. Output to disk requires that you reserve a set of non-SMS disk volumes for use by ABR (see [Section 51.04](#)).

However, if you have an automatic tape library (ATL or silo), you may want to put both copies to tape, since the ATL can quickly mount tapes required for recalls.

In any case, 2 copies are always recommended for safety. Remember that these backups contain the **only** copy of the archived data sets.

- STEP 6** You will need to periodically run a job to execute the REORG function of the ABR utility FDRARCH, described starting in [Section 51.50](#), to purge obsolete entries from the Archive Control File. The frequency of this reorganization will depend on the number of data sets that you archive per day. ABR always adds new entries to the end of the Archive Control File, so it may eventually run out of room for new entries, or become unwieldy. REORG will reorganize the file, deleting entries which are no longer required and compressing all remaining entries, making room for new data. Obsolete entries are usually recognized because they have passed their expiration date, or, for auto-recall entries, they are no longer cataloged for auto-recall.

When REORG detects that the last entry referring to a particular Archive Backup file is being deleted, it will scratch (if on disk) and uncatalog that backup.

- STEP 7** If you archived data sets to disk (STEP 5), you need to delete those disk backup files when they are no longer needed, to make room for new archive data. FDRARCH REORG (STEP 6) will scratch and uncatalog a disk backup if all entries pointing to that backup have been deleted or reached their COPY 1 expiration date.

However, you may want to execute the FDRTSEL utility ([Section 60](#)) to move those COPY 1 Archive Backups on disk to tape when they have expired or are about to expire.

51.32 CONTINUED . . .

STEP 8 For external recalls, described in [Section 51.33](#), an ABR started task is initiated to do the recall by issuing an internal "S SYNRECAL" console command. SYNRECAL is a cataloged procedure (proc) in the system PROCLIB list used by JES2 or JES3. Innovation provides a SYNRECAL proc in the FDR Installation Control Library (ICL) loaded as part of the installation process; you must copy it to a system PROCLIB and make any necessary changes required by your installation. You can change the SYNRECAL name to another; the name is stored in the FDR Global Option Table (ISPF panel A.I.11.2).

STEP 9 If you plan to allow users to request recalls using the remote queue, then the remote queue data set for restore from Archive Backups must be allocated and initialized as discussed in [Section 90](#). You will need to periodically execute an ABR Archive restore job to process this queue and recall the requested data sets (see example in [Section 51.22](#)). Depending on your user's expectations, you may need to schedule this job frequently (e.g., every hour).

STEP 10 In most cases, ABR will recall data sets to the same volume from which they were originally archived. This is acceptable if the original volume is still online and has sufficient free space for the data set. Since data sets might have been archived months or even years ago, their original volumes may no longer exist. Even if they do exist, the volumes may have insufficient space or other errors which will prevent the recall.

You will probably want to enable the ABR RESTORE ALLOCATION LIST, as described in [Section 90](#). This list allows you to specify volume selection rules to be used when a recalled data set can't be allocated on its original volume (it is also used with manual restores from Archive or Volume Backups). When volume serials are deleted from your system, you should update the list to specify alternate output volumes for data sets which were archived from those volumes. You can also use it to pool output volumes, e.g.,

ALLOCATE ALLDSN,VOLG=TSO,NVOLG=TSO

means "for any data set archived from a volume starting with TSO, recall it to any volume starting with TSO".

For SMS-managed data sets, SMS may ignore the volumes specified by the RESTORE ALLOCATION LIST and select the target volume according to SMS rules.

You can ignore the original disk volume serial and use only the RESTORE ALLOCATION LIST to select output volumes; see the LXCONUSE and LXCONVOL options in [Section 51.34](#).

51.33 ABR AUTO-RECALL - OPERATION

ENVIRONMENTS The Catalog Locate Exit supports recalls in a variety of environments: TSO, batch step initiation, dynamic allocation, and CA-ROSCOE. Depending on the environment, one or more recall options are available:

- recall by invoking FDRABR in the address space of the requestor.
- recall by invoking FDRABR as a separate started task but waiting for completion (synchronous external recall).
- recall by invoking FDRABR as a separate started task without waiting for completion (asynchronous external recall).
- recall by adding the request to the Archive Restore remote queue data set.

The Data Set Not Found Exit is not sensitive to the environment; it always processes in the same way, by invoking FDRABR in the address space issuing the OPEN. The recall may fail if there is insufficient free memory in the region to load and execute FDRABR.

INVOKE FDRABR IN THIS ADDRESS SPACE The Catalog Locate Exit or the Data Set Not Found Exit may invoke FDRABR directly to restore the requested data sets in the address space issuing the LOCATE or OPEN. The requesting task will wait until the recall is complete; for TSO users the terminal will be locked while the restore is progressing).

For this type of recall, sufficient storage must be available in the region for ABR to perform the restore. The basic requirement for the ABR restore is 512K, but the requirement may be much more, depending on such factors as whether the backup is compressed by FDR, whether the restore is to a like device or an unlike device, and the DSORG of the data set being restored. If a recall is performed at batch step initiation, the Catalog Locate Exit dynamically increases the REGION size to 1536K if it was smaller than that. In other cases the recall will fail if sufficient storage is not available.

SYNCHRONOUS EXTERNAL RECALL The Catalog Locate Exit initiates a synchronous external recall by issuing a START command for the SYNRECAL cataloged procedure, and then waiting for the SYNRECAL task to complete. The requesting task will wait for the recall to complete; for TSO users the terminal will be locked while the restore is progressing. Since the recall is done in its own address space, there are no special memory requirements.

ASYNCHRONOUS EXTERNAL RECALL The Catalog Locate Exit initiates an asynchronous external recall by issuing a START command for the SYNRECAL cataloged procedure without waiting for SYNRECAL to complete. The requesting task receives control back while the restore is being performed; it will get a "not cataloged" or "not found" indication for the data set

LIMITING CONCURRENT RECALLS In installations where there is a lot of recall activity, you may wish to limit resources used for recalls by limiting the maximum number of concurrent external recalls (started tasks) or total recalls (including those in the address space) . Options LXMAXSTC and LXMAXREC respectively will do this ([see Section 51.34](#)). Recalls which would cause the maximums to be exceeded will wait until other recalls complete. An ABRNOLIM DD statement may be added to a job step to make it exempt from these limits.

REMOTE QUEUE RESTORE The Catalog Locate Exit initiates a remote queue restore by adding a control statement to the Archive Restore remote queue data set, which must have been previously allocated and initialized ([see Section 51.32](#)). The requesting task receives control back; it will get a "not cataloged" or "not found" indication for the data set. The installation must schedule an ABR RESTORE TYPE=ARC job to run at frequent intervals to process the remote queue requests. . If the recall was on behalf of a TSO user, the user receives a terminal message when the restore is completed, after which the recalled data set can be accessed.

51.33 CONTINUED . . .**STEP
INITIATION**

For any batch job, started task, or TSO logon proc, if one or more Job Control DD statements in a step reference data sets that were archived, the Catalog Locate Exit builds a table of those in memory during step initiation. Since step initiation requires a volume serial in order to allocate the data set, the exit will temporarily allocate archived data sets to the volser of an online volume of the proper type.

Just before the job step program (PGM=) is attached, FDRABR is invoked to recall those data sets in the address space of the requestor. If more than one data set is recalled each required tape volume is mounted only once.

If the recall is successful, the exit will reallocate the DDs pointing to the recalled data sets, so that they now point to the volume to which the data set was recalled.

However, recall of several types of data sets cannot be deferred because they are used during other parts of step initiation, or because they cannot be properly reallocated by the Catalog Locate Exit after the recall. For these exceptions, the exit will always invoke an immediate synchronous external recall for each data set as it is encountered:

- STEPLIB and JOBLIB DD statements.
- DCB references (DCB=dsname).
- Data sets that are concatenated to a temporary data set that specifies DISP=(OLD,DELETE).
- Data sets that specify DISP=(OLD,PASS) or DISP=(SHR,PASS) on the DD for the archived data set or for another data set concatenated to the archived data set.
- GDG-all requests. These are DD statements which specify the name of a GDG (Generation Data Group) without including a generation number, e.g., DSN=TEST.GDG, causing MVS to perform an implicit concatenation of all active generations in the GDG.
- Volume references (VOL=REF=dsname), if the referenced data set has been recataloged to MIGRAT. The referenced data set must be recalled immediately because VOL=REF requires the serial number of a real volume, which MIGRAT is not.

**DYNAMIC
ALLOCATION**

If a batch job or started task dynamically allocates an archived data set, ABR will immediately restore the data set by invoking FDRABR in the address space of the requestor. This supports online and database systems such as CICS, DB2, and IMS. You can force a synchronous or asynchronous external recall to be used for a particular job (see the ABRSYNCH and ABRASYNCH DD statements below).

51.33 CONTINUED . . .

TSO If the logon JCL procedure used by a TSO user references archived data sets, they are recalled at step initiation just like any batch job.

All other data sets referenced by TSO, including those used by ISPF, are dynamically allocated by TSO. If a TSO dynamic allocates references an archived data set, an interactive dialog is invoked which informs the user that the data set desired is archived and asks:

- Whether or not to recall the data set. If the user decides not to recall, the dynamic allocation will get a "not found" return code.
 - Whether to recall the data set to the original disk volume from which it was archived, or to a different volume.
 - How to perform the recall: foreground (FG), background (BG), or via the remote queue (RQ).
- If the choice is foreground (FG), the restore will be performed immediately by invoking FDRABR in the TSO user's address space. If the data set has been archived to tape, this requires either that the user have TSO MOUNT authority, or that the installation has set the ALLOCATEFLAG MOUNT option to YES on ISPF panel A.I.4.4.

However, if an ABRSYNCH DD statement is allocated (described later in this section) or the LXFGSYNBN option is set ([See Section 51.34](#)), the restore will instead be performed as a synchronous external recall. Use of an external address space will NOT require the user to have MOUNT authority.

If the recall in the foreground fails for any reason, the LXFGERR option ([see Section 51.34](#)) may direct the Catalog Locate Exit to retry the recall in the background or via the remote queue. This option might be used if you archive data sets with COPY 1 on disk and COPY 2 on tape. If COPY 1 has not yet expired, the foreground recall will recall the data set in a short time, but if COPY 1 is expired and ABR was unable to mount the COPY 2 tape because the user does not have MOUNT authority, it will do the background or remote queue recall.

- If the choice is background (BG), the recall will be performed as an asynchronous external recall. The user will be able to do other work at the terminal while the recall proceeds. The user will be notified when the recall is completed.
- If the choice is remote queue, the recall request will be placed in the remote queue for ARCHIVE restores. The next ABR restore job which references this remote queue will recall the data set. The user will be notified when the restore is completed.

The installation may set options in the FDR Global Option Table which control this dialog ([see Section 51.34](#)). These options can modify any of the bullets above. If you have implemented the ABR RESTORE ALLOCATION LIST you will probably want to suppress the request for a target volume serial. You may want to suppress FG restores, allowing only BG or RQ, or you may want to immediately force one sort of recall (FG, BG, or RQ) without even asking the user.

CA-ROSCOE The Catalog Locate Exit supports Auto-Recall for CA-ROSCOE (Release 5.6 or higher). To activate the support in CA-ROSCOE for Auto-Recall, the installation must install a ZAP from Computer Associates. For example, for CA-ROSCOE release 6.0, the ZAP is GS50978.

Although the CA-ROSCOE support in the Catalog Locate Exit does not support a TSO-style dialog, users can use the ROSCOE command SET DSNRECALL to specify which type of restore to be used if an archived data set is referenced:

- NONE (do not attempt to recall archived data sets, fail request with "not cataloged" error)
- WAIT (synchronous external recall)
- NOWAIT (asynchronous external recall, fail request with "not cataloged" error).

The default is WAIT. For WAIT and NOWAIT, the CA-ROSCOE user will receive messages when the recall is complete.

51.33 CONTINUED . . .

**OPTIONAL
RECALL DD
STATEMENTS**

Depending on the environment, the Catalog Locate Exit or Data Set Not Found Exit may check to see if certain special DDnames have been allocated to the step or TSO user. The presence of these special DDs modifies the operation of Auto-Recall for that step or user. These allocations will specify DUMMY (except for ABRLIB).

In JCL, code the DD statement as

```
//ddname DD DUMMY
```

These are usually used in JCL for batch jobs and started tasks, but you could also include them in TSO logon procs to affect recall options for all users who logon with a given proc.

To allocate these DDs for just one TSO user, the user must go to ISPF Option 6 or TSO READY mode and issue

```
ALLOCATE FILE(ddname) DUMMY
```

**ABRSYNCH
AND
ABRASYNC DD
STATEMENTS**

If an archived data set is recalled because of a dynamic allocation by a batch job or started task, or by a TSO foreground recall request, the Catalog Locate Exit normally recalls the data set by invoking FDRABR within the address space of the requestor. However, if a DDname of ABRSYNCH has been allocated, the restore will instead be performed as a synchronous external recall; if a DDname of ABRASYNC has been allocated, the restore will be performed as an asynchronous external recall (and the dynamic allocation will fail).

Note that an ABRSYNCH or ABRASYNC DD statement has no effect on recalls at step initiation or under CA-ROSCOE.

You can force **all** TSO foreground recalls to be done as synchronous external recalls by setting the LXFGSYNBG option to YES ([See Section 51.34](#)).

ABRIGNOR

If a DDname of ABRIGNOR has been allocated, the Catalog Locate Exit will not recall any data sets for this step or user. The Catalog Locate Exit will normally give a successful (zero) return code to the LOCATE SVC, returning the volume serials currently in the catalog. If the MIGRAT=YES option was used, this may result in an attempt to allocate a disk volume called MIGRAT; if not, and the Data Set Not Found Exit is installed, it may attempt to recall the data set at OPEN time (see ABRIGNRD below).

ABRIGNRD

If a DDname of ABRIGNRD has been allocated, the Data Set Not Found Exit will not recall any data sets for this step or user. If OPEN finds that a data set is not in the VTOC of the allocated volume, it will cause a normal S213-04 ABEND. This can be used to selectively suppress the FDR316 message that is displayed for data sets that were never archived ([see Section 51.31](#)).

ABREMOTE

If a DDname of ABREMOTE has been allocated, the Catalog Locate Exit will place all recall requests due to dynamic allocation by a batch job or started task, or due to TSO foreground recall, into the Archive Restore Remote Queue ([see Sections 51.07](#) and [90.45](#)). The LOCATE will fail so the requesting program will probably generate an error message. The data sets will be recalled later when an ABR Archive Restore which processes the remote queue is executed (see example in [Section 51.22](#)). This might be used with CICS or similar online systems, to accumulate recall requests and minimize tape mounts; the online transaction will fail but will work if executed after the remote queue is processed.

ABRNOLIM

If a DDname of ABRNOLIM has been allocated, the Catalog Locate Exit will make this job step or user exempt from the limits on concurrent recalls set by the LXMAXSTC and LXMAXREC options.

51.33 CONTINUED . . .

ABRLIB If present, the ABRLIB DD statement must point to a program library containing ABR; it can be used to test new versions of ABR. If the Catalog Locate Exit or the Data Set Not Found Exit finds that ABRLIB is present, then the module containing the FDR Global Option Table module (FDROPT) **may** be read from the ABRLIB library; for recalls performed in the address space of the requestor, FDRABR and other modules invoked by FDRABR recall will also be loaded from the ABRLIB library. The following paragraphs explain this in more detail. This is used primarily for testing new versions of ABR ([see Section 90](#)).

Under the Data Set Not Found Exit, ABRLIB is always be used both for FDROPT and FDRABR.

Under the Catalog Locate Exit, for a recall performed in the requestor's address space, ABRLIB is always be used both for FDROPT and FDRABR. Likewise, for a remote queue recall, the remote queue utility (FDRABRUT) is loaded from ABRLIB.

ABRLIB will have no effect on external recalls. The source for the FDRABR and FDROPT modules will depend on the JCL in the SYNRECAL started task. STEPLIB can be used to point to a special FDR program library.

The Catalog Locate Exit accesses FDROPT for various options that affect the operation of auto-recall, described in [Section 51.34](#). These options are set at the time that the Catalog Locate Exit is dynamically installed, so an ABRLIB DD statement cannot override them. However, a test copy of the exit can be installed with a unique set of options; as described in [Section 90](#), the test exit will apply only to certain specified jobs or TSO users. [Section 90](#) also shows how to refresh the copy of FDROPT in use by the exit, to change options dynamically.

51.34 ABR AUTO-RECALL - CUSTOMIZATION OPTIONS

This Section describes customization options for Auto-Recall via the Catalog Locate and Data Set Not Found Exits. They cannot be overridden for individual recalls. The following options affect only the Catalog Locate Exit, unless the description of the option mentions that it affects the Data Set Not Found Exit also.

SETTING OPTIONS The Auto-Recall options are stored in the FDR Global Option Table (FDROPT) and can be changed using the ABR ISPF panels (panels A.I.4.11.2 and A.I.4.11.3, as described in [Section 90](#)). During dynamic installation of the ABR exits the current copy of FDROPT is also loaded into memory; the exits always use that in-memory copy to determine what options are set, so options stored in FDROPT can be changed but will not become active until you relPL or refresh the in-memory copy (see [Section 90](#)).

The topic "Output Volume Selection" at the end of this section discusses the interaction among the various options that affect the chosen target volume, namely LXCONUSE, LXCONVOL, LXCHKSEC, LXEXIT, LXNEWVOL, and LXREISSU.

These options are displayed in the order that they appear on ISPF panels A.I.4.11.2 and A.I.4.11.3.

LXFOREST Specifies whether the Catalog Locate Exit will ask a TSO user to confirm the recall of an archived data set. This option has no effect if LXDFREST and LXUNCAT are both set to NO.

YES – issue message FDRW71 asking the TSO user to confirm or deny the recall. The format of message FDRW71 can be changed by option LXALTMSG

NO – do not issue message FDRW71; the TSO user is not asked to confirm or deny the recall; the recall is forced. If LXDFREST is set to a single restore type, and LXNEWVOL is set to NO, the recall will begin immediately with no user interaction except a message confirming that the recall has started.

The default is YES.

LXFGSYNBG Specifies whether the Catalog Locate Exit is to convert all TSO foreground recall requests to synchronous external recalls.

YES – TSO foreground recalls will be done as synchronous external recalls. This might be useful if TSO users do not have tape mount authority. The user's terminal will be locked until the recall completes. This option can also be invoked for specific users by allocating the ABRSYNCH DDname (see [Section 51.33](#)).

NO – TSO foreground recalls will be executed in the address space of the TSO user.

LXDFREST Specifies the restore types that are available for Auto-Recalls by TSO users under the Catalog Locate Exit. The values that may be specified are NO or any combination of FG, BG, and RQ. If more than one restore type is specified, the TSO user will be prompted to choose which restore type to use. If only one restore type is specified, that restore type will be used automatically, and the TSO user will not be prompted. The restore types are described in [Section 51.33](#).

NO – Auto-Recall is not available for TSO users.

FG – foreground recalls are available.

BG – background recalls are available.

RQ – remote queue recalls are available.

The default is FG,BG,RQ; that is, all of the restore types are available, and the user will be prompted to choose one.

51.34 CONTINUED . . .

LXFGERR Specifies whether a recall should be retried in the background or via the remote queue, if it fails in the TSO foreground. This option has no effect unless option LXDFREST includes FG as an available restore type.

NO – a failed foreground recall will not be retried.

BG – a failed foreground recall will be retried as a background recall.

RQ – a failed foreground recall will be retried as a remote queue recall.

The default is NO.

LXFGERR is an important option in the environment where data sets are archived to disk with a short retention and to tape with a longer retention, and the TSO users do not have MOUNT privileges. As long as an archived data set is available on disk, it can be quickly be recalled in the foreground. After the disk copy expires, ABR will automatically try to restore from the tape copy, but the dynamic allocation will fail because the user does not have MOUNT privileges. If LXFGERR is set to BG or RQ, the recall will automatically be sent to the background or the remote queue, where the tape can be mounted.

Note that you can specify BG or RQ even if BG or RQ is not specified as a valid restore type under LXDFREST. For example, if you specify FG as the only value for LXDFREST, and you specify BG as the value for LXFGERR, then all recalls will initially be tried in the foreground, and those that fail will automatically be retried as an asynchronous external recall.

LXNOMSG Specifies whether the Catalog Locate Exit will issue or suppress messages FDRW70 and FDRW79, when the exit finds that the user has referenced a data set that is archived for recall, but the installation does not allow Auto-Recalls for TSO users. This option has no effect unless LXDFREST and LXUNCAT are both set to NO.

YES – suppress the messages. The user will not be informed that the data set was archived.

NO – issue messages FDRW70 and FDRW79 to inform the user that the data set was archived, in case they want to restore it manually.

The default is NO.

LXNEWVOL Specifies whether the Catalog Locate Exit will permit a TSO user to designate a new output volume to which to restore an archived data set.

YES – issue messages FDRW76 and FDRW77, telling the TSO user the volume that is designated as the output volume for the restore, and permitting a different volume to be specified.

NO – messages FDRW76 and FDRW77 will not be issued, and the TSO user will not have the option to override the chosen volume serial. You may want to set NO if you use the ABR RESTORE ALLOCATION LIST to direct restores to volumes, or if most recalled data sets are SMS-managed.

The default is YES.

LXREISSU Specifies whether the Catalog Locate Exit will re-display the designated volume and give the TSO user another chance to change it, after the user has designated a new output volume to which to restore an archived data set (see LXNEWVOL above). This option has no effect if either LXDFREST or LXNEWVOL is set to NO.

YES – when the user has replied to message FDRW77 by designating a new output volume to which to restore the data set, the Catalog Locate Exit will reissue messages FDRW76 and FDRW77. Message FDRW76 will re-display the new volume that the user has designated, and message FDRW77 will give the user another chance to change it.

NO – the TSO user will have only one chance to designate a new output volume to which to restore the data set.

The default is NO.

51.34 CONTINUED . . .

LXSYNPROC Specifies the name of the cataloged procedure (proc) that will be invoked by the Catalog Locate Exit to perform external recalls as described in [Section 51.31](#). The exit will issue an internal "START procname" to start the recall task, passing parameters identifying the data sets to be recalled.

The default is SYNRECAL.

Note: When installing the Catalog Locate Exit, you must copy the SYNRECAL cataloged procedure from the ABR Installation Control Library (ICL) to a cataloged procedure library that is available to JES for START commands, customizing it if necessary. If you rename it, specify the new name for this option.

LXCHKSEC Specifies whether the Catalog Locate Exit itself will perform security checking before recalling an archived data set. This is in addition to security checking that may be done by FDRABR during the restore or by FDRABRUT during the remote queue update. [Section 51.35](#) "Security Considerations" gives further details. This option has no effect if global options ALLCALL and SECEXIT are all set to NO.

YES – the exit will perform security checking. In addition, if the LXCHKSEC and SECEXIT options are both set to YES, the FDR Data Set Security exit may designate a new volume as the output volume for the restore, also as detailed in [Section 51.35](#).

NO – the Catalog Locate Exit will not perform security checking but FDRABR or FDRABRUT may do so.

The default is NO.

LXBYPSEC Specifies whether the Catalog Locate Exit and the Data Set Not Found Exit will cause security checking to be bypassed during the recall. This option is separate from the LXCHKSEC and LXEXIT options, which control security checking that may be done by the Catalog Locate Exit itself. [Section 51.35](#) "Security Considerations" gives further details.

YES – inhibit all security checking done by FDRABRUT to add the restore request to the remote queue (Locate Exit only), or by FDRABR to perform the recall. The bypass is effective for RACF, CA-TOP SECRET, CA-ACF2, and other security systems. The bypass applies both to security checking that FDRABR and FDRABRUT perform explicitly, and to security checking performed by Operating System routines that FDRABR and FDRABRUT invoke.

NO – normal security checking will be performed by FDRABR and FDRABRUT. This may mean that users will not be able to recall data sets unless they have the authority to allocate them.

The default is NO.

LXEXIT Specifies whether the Catalog Locate Exit will invoke a special security exit. [Section 51.35](#) "Security Considerations" gives further details.

YES – a special Auto-Recall security exit will be called, detailed in [Section 51.35](#). The name of the exit must be specified in the LXEXIT ... EXIT NAME option. In addition to checking security, the LXEXIT exit may designate a new volume as the output volume for the restore.

NO – no special security exit is called.

The default is NO. There is no default for the name of the exit; if you activate the exit you must specify its name.

LXNCDENY Specifies whether the Catalog Locate Exit will give a 'not cataloged' return code to LOCATE, when a TSO user references a data set that has been archived for recall, and specifies that the data set should not be recalled. This option has no effect if options LXDFREST and LXUNCAT are both set to NO.

YES – LOCATE will receive a 'not cataloged' return code when a TSO user denies the recall request.

NO – LOCATE will receive a normal (zero) return code when a TSO user denies the recall request.

The default is YES.

51.34 CONTINUED . . .

LXCONUSE Specifies whether the Catalog Locate Exit will initially designate the constant new volume serial specified by the LXCONVOL option as the output volume for Auto-Recall.

YES – the value of LXCONVOL will be passed as the output volume serial for all recalls.

NO – the original volume serial (or MIGRAT) from the Auto-Recall catalog entry is passed.

The default is NO.

The usual use of the LXCONUSE/LXCONVOL options is to designate a non-existent volume in order to force ABR to select a volume using only the RESTORE ALLOCATION LIST, which must be enabled for non-SMS data sets. You may want to set LXNEWVOL to NO so that users cannot override the constant volume.

LXCONVOL Specifies the constant volume serial that the Catalog Locate Exit will specify as the output volume for Auto-Recall, if LXCONUSE is set to YES.

The default is NEWVOL.

LXALTMSG Specifies whether the Catalog Locate Exit will use an alternate format for message FDRW71, when asking a TSO user whether an archived data set should be recalled, so as to require a positive action (keying in 'YES' instead of just pressing 'Enter') to cause an archived data set to be recalled. This option has no effect if options LXDFREST and LXUNCAT are both set to NO, or if option LXFOREST is set to NO.

YES – the message will be:

FDRW71 TYPE 'YES' TO PERMIT THE RESTORE OR PRESS 'ENTER' TO BYPASS

NO – the message will be:

FDRW71 TYPE 'END' TO BYPASS THE RESTORE OR PRESS 'ENTER' TO CONTINUE

The default is NO.

LXUNCAT Specifies whether the Catalog Locate Exit will offer a TSO user the choice of uncataloging an archived data set instead of recalling it. If option LXFOREST is set to NO, then option LXUNCAT will be ignored.

YES – it will offer the choice of uncataloging the data set. If the user chooses to uncatalog, the data set will no longer be eligible for Auto-Recall by the Catalog Locate Exit. The LOCATE SVC will get a "not found" return code.

NO – it will not offer the choice of uncataloging the data set.

The default is NO.

LXMAXSTC
LXMAXREC Specifies the maximum number of started tasks for external recalls (LXMAXSTC) or total concurrent recalls (LXMAXREC) that the Catalog Locate Exit will allow to be running in the system at one time. Each may be from 1 to 255. If both are specified, then LXMAXSTC must be less than or equal to LXMAXREC.

If either limit is reached, then:

- If a batch job is requesting the recall, the Catalog Locate Exit will issue a message to the console operator, and will go into a wait for one minute at a time, until the number of started tasks or total recalls falls below the limit. The console operator can cancel the job while it is waiting.

- If a TSO user is requesting the recall, the Catalog Locate Exit will issue a message to the TSO terminal, asking the user whether he wants to wait, or to terminate the current operation and perhaps try again later. If the TSO user chooses to wait, then the Catalog Locate Exit will go into a wait for one minute at a time, until the number of started tasks or total recalls falls below the limit. The TSO user can use the Attention (PA1) key to terminate the wait if he changes his mind.

The default (NONE) is that there is no limit on the number of started tasks for recall or total recalls that the Catalog Locate Exit will allow to be running in the system at one time.

51.34 CONTINUED . . .

LXDIRVOL When ABR Auto-Recall is used under JES3, the Catalog Locate Exit is invoked during batch job scheduling, but it may not be on the system where the job will eventually run, so the recall cannot be done at that time; data sets will be recalled when the job actually runs. The exit will return a disk volume serial so that JES3 will think the data set is on that volume, but if the volume returned is not online to all systems in the JES3 complex, JES3 may not be able to schedule the job on the proper system, or may fail it. For JES3 users only, these options identify the volume serial (LXDIRVOL) and device type such as 3390 (LXDIRTYP) to be passed during JES3 scheduling. It must be a disk volume online to all systems in the JES3 complex.

OUTPUT VOLUME SELECTION Whenever the Catalog Locate Exit initiates a recall of an archived data set, using any restore type, it designates a target output volume to which the data set should be restored. The default is that the designated volume is the volume to which the data set is cataloged, which is either the volume from which the data set was archived or 'MIGRAT' if the MIGRAT option was in effect when the data set was archived.

The following explains the interaction among the various options that may change the designated output volume. These options are LXCONUSE, LXCONVOL, LXCHKSEC, LXEXIT, LXNEWVOL, and LXREISSU.

- First, the exit sets the default of the volume to which the data set is cataloged.
- But, If the LXCONUSE option is set to YES, the volume specified by the LXCONVOL option is set.
- If the LXEXIT option is set to YES, the Catalog Locate Exit calls the user-written Auto-Recall exit named by the LXEXIT ... EXIT NAME option. The LXEXIT exit may designate a new volume as the output volume for the restore, as discussed in the [Section 51.35 "Security Considerations"](#).
- If the LXCHKSEC and SECEXIT options are both set to YES, the Catalog Locate Exit calls the FDR Data Set Security exit, using the name specified by the SECEXIT ... EXIT NAME option. The Data Set Security exit may designate a new volume as the output volume for the restore, as discussed in the [Section 51.35 "Security Considerations"](#).
- If the archived data set was referenced by a TSO user, and the LXNEWVOL option is set to YES, the Catalog Locate Exit issues messages FDRW76 and FDRW77, telling the TSO user the volume that is currently designated as the output volume for the restore, and allowing the TSO user to designate a new volume. If the LXREISSU option is also set to YES, the exit reissues messages FDRW76 and FDRW77, re-displaying the new volume that the just entered, and allowing the user to change the designated volume again; this continues until he hits ENTER without changing the volume again.

It would be very unusual to have all of the above options set to YES at the same installation. For example, if an installation uses an LXEXIT exit or a Data Set Security exit (LXCHKSEC and SECEXIT options) to designate the output volume for the restore, that installation would probably want to set the LXNEWVOL option to NO so that the TSO user could not override the designated volume.

NOTE: the volume that the Catalog Locate Exit designates as the output volume for the restore is not necessarily the volume to which the data set will actually be restored. If the designated volume is 'MIGRAT', ABR will substitute the original volume from which the data set was archived, as recorded in the Archive Control File. If the designated volume is not online or is full (or any other allocation failure), ABR will use the RESTORE ALLOCATION LIST ([Section 90](#)), if enabled, to select a different volume. LXCONUSE and LXCONVOL, in particular, are often used to designate a non-existent volume in order to force ABR to select a volume using the RESTORE Allocation Control List. Also, if SMS is active, SMS may assign the data set to a different volume ([Section 70](#)).

For multi-volume data sets, the above rules only affect the first volume of the data set. Additional pieces of the data set will get a target volume of the original volume serial. Only the RESTORE ALLOCATION LIST or SMS can override that volume.

51.35 ABR AUTO-RECALL - SECURITY

There are special security considerations for auto-recall of archived data sets.

As described in [Section 51.03](#), allocating and restoring a data set requires ALTER authority to that data set. Even adding a restore request to a remote queue with FDRABRUT will require ALTER authority. But recalling a data set is not really creating it; it is simply moving it back to disk from a less expensive storage medium such as tape. A user with only READ authority to a data set should be able to recall a data set and open it for input.

This dilemma can be avoided by activating the LXCHKSEC and LXBYPSEC options ([see Section 51.34](#)).

LXCHKSEC When LXCHKSEC is YES, the Catalog Locate Exit will check the caller's authority to the data set (in security class DATASET) before it even invokes FDRABR to do the recall or FDRABRUT to add the request to the remote queue. If the user does not have the proper authority, the recall will not be attempted.

It will check for READ authority to the data set, so that a user with only read access can recall the data set and read it. This does not present a security exposure even if the user is trying to write to the data set, because after the data set is recalled, OPEN will still perform the normal security checks. At worst, some disk space will be wasted by recalling a data set that the user will not be able to access, until the data set becomes eligible to be archived again.

However, when the output volume has been changed (by the LXCONUSE option, the LXEXIT recall exit, or the FDR Data Set Security Exit), the security check is for ALTER authority. This does not apply if a TSO user modifies the output volser via a response to the FDRW77 message.

This security check is always done in the address space of the requestor, even if the recall will be performed as an external recall.

The LXCHKSEC option will not check for authority to the output volume (in the DASDVOL security class) because the output volume serial may be changed by ABR or SMS.

The Data Set Not Found Exit does no security checking of its own.

LXBYPSEC When LXBYPSEC is YES, the Catalog Locate Exit and Data Set Not Found Exit will cause all security checking to be bypassed when they invoke FDRABRUT to add the restore request to the remote queue (Locate Exit only), or FDRABR to perform the restore. All security checks from FDRABR, FDRABRUT, and system services (such as allocation) will be suppressed. This allows any user to recall any archived data set, but there is no security exposure since normal security checking will be back in effect when the data set is OPENed.

This also allows ABR to read Archive Backup files to which the user may not be authorized at all.

LXBYPSEC will bypass security checking for all foreground and external recalls. But for remote queue restores, it bypasses the security checks done by FDRABRUT before adding the request to the queue, but does not affect the actual restore. The ABR RESTORE TYPE=ARC job which processes the request on the remote queue must have authority to allocate and restore all data sets on all volumes. You may want to run that restore job under a userid that is exempt from security checking (e.g. RACF OPERATIONS attribute).

The security bypass invoked by LXBYPSEC is effective for RACF, CA-ACF2 and CA-TOP SECRET.

RECOMMENDATION Innovation recommends that you enable both LXCHKSEC and LXBYPSEC. This allows users with only READ authority to recall a data set.

51.36 ABR AUTO-RECALL - HSM COMPATIBILITY

ABR Auto-Recall is designed to be compatible with DFHSM (the IBM HSM product) and DFSMSHsm (the HSM component of the IBM DFSMS/MVS product and OS/390). Also, if both ABR and DFHSM or DFSMSHsm are active on a system, ABR will coexist with HSM, with each product performing recalls for only the data sets which have been archived/migrated by it.

**HSM
COMPATI-
BILITY**

If the ABR option MIGRAT=YES is used for Archive Backups, then data sets archived for recall will be recataloged to a volser of MIGRAT, just as HSM does, as well as having the normal ABR ARCHIVE indicators in the catalog entry. The ABR Catalog Locate Exit will properly recall data sets cataloged to MIGRAT in most cases, but some software products test for a volser of MIGRAT and issue a direct HSM request for the recall (the HRECALL command, or the HSM ARCHRCAL macro, or the HSM SVC); such products include:

DB2

NFSS (Network File System Server)

JES2 (for archived JCLLIBs (user proclibs))

and others.

ABR provides two modules to handle these requests:

- FDR00024, a front-end to the HSM SVC. When it is entered, it tests to see if the requested data set was really archived by ABR and invokes ABR to recall them. Other requests are passed to HSM if it is active.
- FDRGIVER, an alternative to the IBM ARCGIVER module. ARCGIVER is an HSM module invoked by the ARCHRCAL macro and other HSM macros, when invoked it issues the HSM SVC to request HSM services. FDRGIVER does the same except that it only handles recall and delete requests, allowing ABR to handle HSM requests issued by other programs.

If the ABR Catalog Locate Exit is being dynamically installed (as described in [Section 90](#)), ABR will also dynamically install the HSM SVC front-end, FDR00024. If there is no other module called ARCGIVER in the system, ABR will also install FDRGIVER in place of ARCGIVER. With these modules in place, archived data sets will be recalled even for programs which issue direct HSM requests.

These modules also handle HSM delete requests (e.g., the HDELETE command) for ABR archived data sets by uncataloging them.

Other HSM requests (to migrate, restore, etc.) will be passed to HSM, if it is active.

**HSM
COEXISTENCE**

The HSM SVC front-end, plus code in the ABR Catalog Locate Exit, allows ABR to coexist with DFHSM or DFSMSHsm when both products are active, even when the ABR MIGRATE=YES option is used (so that both HSM and ABR catalog archived and migrated data sets to volser MIGRAT).

When HSM is installed (even if it is not active), the HSM version of the ARCGIVER module will be used. Since all requests to ARCGIVER are translated into the HSM SVC, so the HSM SVC front-end will identify which system archived a given data set and pass the recall or delete request to ABR or HSM.

The result is that data sets archived by ABR will be recalled by ABR, and those migrated by HSM will be recalled by HSM.

51.40 REMOTE QUEUE UTILITY (FDRABRUT)

Your installation may decide to allow end-users to request ABR processing of the data sets they own. Although it is possible for a user to submit his own ABR job, it is wasteful to create an incremental backup or Archive Backup file for just a few data sets. Restores can also be made more efficient if the requests are "batched" together, minimizing tape mounts.

As noted in [Section 50](#) and elsewhere in [Section 51](#), ABR supports optional remote queues. These queues allow users to request certain ABR functions to be performed on data sets they specify. The remote queues are not processed immediately. The requests are handled when an appropriate ABR batch job is run by Data Center Operations or a scheduling program.

The remote queues are managed by the utility program FDRABRUT and can be used for 4 types of requests:

- include a data set in ABR Volume Backups
- archive a data set
- restore a data set from Volume Backups
- restore a data set from Archive Backups

You can support all 4 types of requests, or any subset of them.

For all 4 request types, a remote queue data set can be used. These are simple sequential data sets, and control statements specifying the user's request are added to the end of them. There is a separate remote queue data set for each type of request. Users can use one of several techniques to add requests to a queue (all involving FDRABRUT) and an appropriate ABR job (backup, archive, restore) will read and process the requests. Since it is possible that users are adding requests at the same time that ABR is reading them, an internal ENQ is used to serialize access to each remote queue data set. ABR empties the data set when the requests are processed.

For the first 2 request types (backup and archive), you can optionally set a flag in the Format 1 DSCB of the data set to be processed. For Volume Backup, it simply sets the update flag which will cause it to be included in the next incremental backup. For Archive, it sets a flag in the reserved byte used by ABR, causing it to be archived. These are effective only if their disk volume is processed in the next appropriate ABR run. This option is used if DISKUPDATE=YES is set in the FDR Global Option Table ([see Section 90](#)) but it will work only if FDRABRUT can be executed as an authorized program from an authorized library.

TSO USERS TSO users have several ways of executing FDRABRUT to add requests to a remote queue:

- they can execute it directly, e.g.,
`CALL LIB(FDRABRUT)`
 and interactively enter the control statements documented in the following sections.
- many ABR ISPF panels have a "remote queue" option, to add the request to a queue. FDRABRUT statements are generated by the panel and the utility is dynamically invoked.
- there is an ABF ISPF panel devoted to remote queue operations of all types. FDRABRUT statements are generated by the panel and the utility is dynamically invoked.
- when a TSO user references an archive data set, one of the options they may be presented is to add the request to a queue (RQ). FDRABRUT statements are generated by the ABR Catalog Locate Exit and the utility is dynamically invoked.

If the request is for backup or archive, and the DISKUPDATE=YES option is set in the FDR Global Option Table, FDRABRUT will fail unless it is executed as an authorized program under TSO. See [Section 90](#) for information on authorizing programs under TSO.

51.40 CONTINUED . . .

BATCH JOBS Batch jobs can execute FDRABRUT directly, using the JCL and control statements documented in the following sections. This might be used to request that a newly created data set be archived almost immediately, by putting an FDRABRUT step in the creating job.

SECURITY If the ALLCALL option is enabled in the FDR Global Option Table (ISPF panel A.I.1, [see Section 90](#)), FDR security checking is enabled and users will only be able to add requests to the various remote queues if they have a proper level of authority to each data set being added. FDRABRUT will check for the proper authority before adding the request.

If DISKUPDATE=YES is in effect for backup requests, FDRABRUT will follow the rules defined in [Section 50.03](#) (for backup) and 51.03 (for archive). However, if DISKUPDATE=NO, FDRABRUT will not do a volume authority check; it will only check for authority to the individual data sets.

For backup requests when DISKUPDATE=NO is in effect, and for all restore requests, the use of DSG= to select groups of data sets will be disabled when security is enabled, since FDRABRUT will not be able to determine the names of the individual data sets for security checks. Only DSN= with a fully-qualified data set name is supported.

THE REMOTE QUEUE DATA SETS [Section 90](#) has details on creating and identifying the 4 remote queue data sets. You may choose to implement any or all of them.

FDRABRUT will dynamically allocate the proper remote queue data set for each control statement, so the data set names of the remote queues must be specified in the FDR Global Option Table (panel A.I.4.5). There is a default name for each data set, but you can change it to anything you prefer. [Section 90](#) shows the default names.

SERVICING THE REMOTE QUEUES An appropriate ABR job must be run periodically to process any requests on each of the queues. For the backup and archive queues, they can be included in your regular daily Volume Backup or Archive Backup ABR jobs. The requested data sets are simply included with those that are selected for other reasons.

For the two restore queues, you will need to execute a periodic job of the proper type (RESTORE TYPE=ABR or TYPE=ARC) which will restore only the selected data sets. In order to meet your user's expectations, you should run this job at regular intervals (e.g., once an hour, once every 4 hours).

If DISKUPDATE=YES is in effect for the backup and archive queues, data sets on the queue will be processed the next time that an ABR backup or archive job processes the volumes containing those data sets. Data sets on the backup queue are automatically included in the next incremental, even if they have not been updated, while data sets on the Archive queue may not be archived unless you specify the REMOTE operand on the DUMP TYPE=ARC statement.

For the two restore queues, and also the backup and archive queues if DISKUPDATE=NO is in effect, use of the remote queue for a particular ABR function is controlled by the presence of a particular DDNAME in the JCL of the ABR step. Here are the DDnames and the functions that invoke them:

ABRBKDQ	DUMP TYPE=ABR/AUTO/DSF
ABRARDQ	DUMP TYPE=ARC
ABRREST	RESTORE TYPE=ABR
ABRARCH	RESTORE TYPE=ARC

For more details, see the JCL Requirements for backup and restore in [Sections 50](#) and [51](#).

51.41 FDRABRUT JOB CONTROL REQUIREMENTS

The following Job Control Statements are required to execute FDRABRUT:

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRABRUT resides. The library must be APF authorized if DISKUPDATE=YES is set in the FDR Global Option Table.

EXEC STATEMENT Must specify the program name (PGM=FDRABRUT) and region requirement (REGION=256K or more).

SYSPRINT DD STATEMENT Specifies the output message data set; it is required. It is usually a SYSOUT data set. DCB characteristics are RECFM=FBA and LRECL=120; the blocksize will default to 1210 on disk or tape.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

//ABNLIGNR DD DUMMY

ABRARCH DD STATEMENT Optionally specifies the remote queue data set for Archive restore requests. The data set name must contain the character string 'ABRARCH' in other than the first index level. If omitted, the archive restore remote queue data set whose name is in the FDR Global Option Table will be dynamically allocated.

ABRREST DD STATEMENT Optionally specifies the remote queue data set for Volume Backup restore requests. The data set name must contain the character string 'ABRREST' in other than the first index level. If omitted, the backup restore remote queue data set whose name is in the FDR Global Option Table will be dynamically allocated.

ABRARDQ DD STATEMENT Optionally specifies the remote queue data set for Archive Backup requests. The data set name must contain the character string 'ABRARDQ' in other than the first index level. If omitted, the Archive Backup remote queue data set whose name is in the FDR Global Option Table will be dynamically allocated.

ABRBKDQ DD STATEMENT Optionally specifies the remote queue data set for Volume Backup requests. The data set name must contain the character string 'ABRBKDQ' in other than the first index level. If omitted, the Volume Backup remote queue data set whose name is in the FDR Global Option Table will be dynamically allocated.

SYSIN DD STATEMENT Specifies the control statement data set. Usually a DD * data set.

51.42 FDRABRUT ARCHIVE/BACKUP STATEMENT

ARCHIVE **DSG=**dsgroup
BACKUP **DSN=**dsname

 ,**ADATE=**yyddd|yyyyddd
 ,**ADAYS=**nnnn

 ,**NOTIFY=**userid

 ,**VOL=**vvvvvv
 ,**VOLG=**vvvvvv

ARCHIVE STATEMENT The ARCHIVE statement requests that either specific data sets or groups of data sets be archived. It will either mark the DSCB of the data set(s) to be archived (DISKUPDATE=YES) or store the archive request in the Archive Backup remote queue data set (DISKUPDATE=NO).

BACKUP STATEMENT The BACKUP statement requests that either specific data sets or groups of data sets be included in the next incremental backup. It will either mark the DSCB of the data set(s) to be backed up (DISKUPDATE=YES) or store the backup request in the Volume Backup remote queue data set (DISKUPDATE=NO).

OPERANDS **DSN=** Specifies the fully-qualified data set name, up to 44 characters, for a data set to be marked for later archiving or backup. If neither VOL= nor VOLG= are specified, a catalog LOCATE will be done to get the volume serials associated with this name.

DSG= Specifies the prefix of a group of data sets to be marked for later archiving or backup, up to 44 characters. The VTOC of every volume specified by the VOL= or VOLG= operands will be scanned for data sets beginning with this string. If DISKUPDATE=NO is in effect and FDR security is enabled, use of DSG= is disabled (see "Security" in [Section 51.40](#)).

 There is a special form of the DSG= operand. Leading periods after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one index level is to be bypassed.

Note: either DSN= or DSG= is required, but not both.

ADATE= Only used with ARCHIVE, specifies a Julian date which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). 2-digit year numbers less than 70 are assumed to be 20xx. Data sets will not be archived unless the last reference date stored by IBM in the Format 1 DSCB is less than this value.

ADAYS= Only used with ARCHIVE, specifies a number of days. Data sets will not be archived unless the last reference date stored by IBM in the Format 1 DSCB is less than today's date minus this value.

Note: either ADATE= or ADAYS= can be used, not both.

NOTIFY= Specifies a TSO userid to be notified when the data set has been archived or backed up. This is ignored if DISKUPDATE=YES is in effect; it works only when the remote queue data set is used.

VOL= Specifies the volume serial of the disk volume on which to search for the data set or data set group.

51.42 CONTINUED . . .

VOLG= Specifies a volume serial prefix. All online volumes whose serial starts with the prefix will be searched for the data set or data set group.

Note: either VOL= or VOLG= must be present, except that if DSN= names a cataloged data set, only the volume from the catalog will be searched.

51.43 FDRABRUT RESTORE STATEMENT

RESTORE	ARCHIVE	,NEWN=dsname
	BACKUP	,NEWG=newgroup
		,NEWI=newindex
	,DSN=dsname	
	,DSG=dsgroup	,NOCAT
		,RECAT
	,ADATE=yyddd	
		,NOTIFY=userid
	,BLKF=nn	
		,NVOL=newvolser
	,COPY=n	
		,RLSE
	,DATA=ALL	,%FREE=nn
	,DATACLAS=dataclass	,STORCLAS=storageclass
	,NULLDATACLAS	,NULLSTORCLAS
	,DSNENQ=NONE	,TRK=nnn
		,CYL=nnn
	,MGMTCLAS=managementclass	
	,NULLMGMTCLAS	,VOL=vvvvvv
		,VRECAT

RESTORE ARCHIVE STATEMENT The RESTORE ARCHIVE statement requests that either specific data sets or groups of data sets be restored from Archive Backup. It will store the restore request in the Archive restore remote queue data set. When the ABR job processing this remote queue is run, it will search the Archive Control File for the backups of the data sets specified, according to the rules in [Section 51.02](#) and restore them to disk.

RESTORE BACKUP STATEMENT The RESTORE BACKUP statement requests that either specific data sets or groups of data sets be restored from Volume Backup. It will store the restore request in the Volume Backup restore remote queue data set. When the ABR job processing this remote queue is run, it will search the for the backups of the data sets specified, according to the rules in [Section 50.02](#) and restore them to disk.

OPERANDS Either ARCHIVE or BACKUP is required. For a description of all other operands see [Section 51.09](#) for RESTORE ARCHIVE and 50.09 for RESTORE BACKUP. Note that ADATE= only applies to RESTORE ARCHIVE, while GEN= and CYCLE= only apply to RESTORE BACKUP.

If FDR security is enabled, use of DSG= is disabled (see "Security" in [Section 51.40](#)).

51.44 FDRABRUT RESET/REMOVE STATEMENT

	RESET	ARCHIVE
	REMOVE	BACKUP
		,DSN=dsname ,DSG=dsgroup ,VOL=vvvvvv ,VOLG=vvvvvv
RESET STATEMENT	The RESET statement is used to remove requests currently queued for backup or archive. If DISKUPDATE=YES is in effect, it will reset the appropriate flag in the Format 1 DSCB of the data set; for DISKUPDATE=NO it will search the appropriate remote queue data set and remove the requested entries.	
REMOVE STATEMENT	The REMOVE statement is used to remove requests currently queued for restore. It will search the appropriate remote queue data set and remove the requested entries.	
OPERANDS	ARCHIVE	remove requests from the Archive Backup or restore queue.
	BACKUP	remove requests from the Volume Backup or restore queue. Note: Either ARCHIVE or BACKUP must be specified, not both.
	DSN=	Specifies the fully-qualified data set name, up to 44 characters, for a data set to be listed.
	DSG=	Specifies the prefix of a group of data sets to be listed, up to 44 characters. If FDR security is enabled, use of DSG= is disabled (see "Security" in Section 51.40). There is a special form of the DSG= operand. Leading periods after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one index level is to be bypassed. Note: either DSN= or DSG= can be used, not both. If neither DSN= nor DSG= is specified, all data sets on the queue will be removed if they match the VOL= or VOLG= operand
	VOL=	Specifies the volume serial of the disk volume for which the request was made.
	VOLG=	Specifies the volume serial prefix of the volumes for which the request was made. Note: either VOL= or VOLG= can be used, not both. If both VOL= and VOLG= are omitted, then data sets are removed from the queue if they match the DSN= or DSG= operands.

51.45 FDRABRUT LIST STATEMENT

LIST

**ARCHIVE
BACKUP**

**,DUMP
,RESTORE**

,DSN=dsname
,DSG=dsgroup

,VOL=vvvvvv
,VOLG=vvvvv

LIST STATEMENT The LIST statement is used to list requests currently queued on one of the remote queue data sets. For backup and archive requests, if DISKUPDATE=YES was in effect, it will not list those requests; it will only list those on a remote queue data set (DISKUPDATE=NO).

OPERANDS

ARCHIVE List requests from the Archive Backup or restore queue.

BACKUP List requests from the Volume Backup or restore queue.
Note: Either ARCHIVE or BACKUP must be specified, not both.

DUMP List requests from the Archive or Volume Backup queue.

RESTORE List requests from the Archive or Volume restore queue.
If neither DUMP or RESTORE is specified, RESTORE is assumed.

DSN= Specifies the fully-qualified data set name, up to 44 characters, for a data set to be listed.

DSG= Specifies the prefix of a group of data sets to be listed, up to 44 characters.
There is a special form of the DSG= operand. Leading periods after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one index level is to be bypassed.
Note: either DSN= or DSG= can be used, not both. If neither is specified, all queued data sets will be listed.

VOL= Specifies the volume serial of the disk volume for which the request was made.

VOLG= Specifies the volume serial prefix of the volumes for which the request was made.
Note: either VOL= or VOLG= can be used, not both. If both are omitted, then requests are listed no matter what disk volser is associated with them.

51.46 FDRABRUT EXAMPLES

The following examples illustrate some of the ways of executing the remote queue utility in batch for backup and restore of data sets.

EXAMPLE 1 Queue requests to include 2 data sets in Volume Backups, and to archive 2 other data sets. The data set OLD.DATA is not cataloged, so its volume is identified; other data sets will be located in the system catalogs. If the DISKUPDATE=YES option is specified in the FDR Global Option Table, FDRABRUT will mark the data set's DSCB; if not, it will allocate the Volume Backup and Archive Backup remote queue data sets and add the requests to them.

```
//BACKUP      EXEC      PGM=FDRABRUT
//SYSPRINT    DD        SYSOUT=*
//SYSUDUMP    DD        SYSOUT=*
//SYSIN       DD        *
      BACKUP    DSN=XYZ.PAYROLL
      BACKUP    DSN=USERID
      ARCHIVE   DSN=PAYROLL.A
      ARCHIVE   DSN=OLD.DATA,VOL=PAY005
```

EXAMPLE 2 Queue requests to restore two data sets from Archive Backup and one from Volume Backup. Data set PAYROLL.ABC will be restored as PAYROLLX.ABC, while PAYROLL.AXY will be restored to volume PROD92. DD statements are provided to point to the Archive Backup and Volume Backup remote queue data sets.

```
//RESTORE     EXEC      PGM=FDRABRUT
//SYSPRINT    DD        SYSOUT=*
//SYSUDUMP    DD        SYSOUT=*
//ABRARCH     DD        DSN=FDRABR.ABRARCH.SYSTEM1,DISP=SHR
//ABRREST     DD        DSN=FDRABR.ABRREST.SYSTEM1,DISP=SHR
//SYSIN       DD        *
      RESTORE   ARCHIVE,DSN=PAYROLL.ABC,NEWI=PAYROLLX
      RESTORE   ARCHIVE,DSN=PAYROLL.AXY,NVOL=PROD92
      RESTORE   BACKUP,DSN=SYS2.LIBRARY
```

EXAMPLE 3 Remove previously queued backup and restore requests, from the Archive queues and the Volume Backup queues. The request for SYSX.LIBRARY is removed only if it is marked as coming from volume TSOXXX.

```
//REMOVE      EXEC      PGM=FDRABRUT
//SYSPRINT    DD        SYSOUT=*
//SYSUDUMP    DD        SYSOUT=*
//SYSIN       DD        *
      REMOVE    ARCHIVE,DSN=USER.FILE
      REMOVE    BACKUP,DSN=SYSX.LIBRARY,VOL=TSOXX
      RESET     ARCHIVE,DSN=USER.FILE
      RESET     BACKUP,DSN=SYSX.LIBRARY,VOL=TSOXXX
```

EXAMPLE 4 List all previously queued requests for the Archive restore queue and the Volume Backup queue.

```
//LIST        EXEC      PGM=FDRABRUT
//SYSPRINT    DD        SYSOUT=*
//SYSUDUMP    DD        SYSOUT=*
//SYSIN       DD        *
      LIST      ARCHIVE
      LIST      BACKUP,DUMP
```

51.50 ARCHIVE MAINTENANCE UTILITY (FDRARCH)

FDRARCH is an ABR utility which performs required maintenance on the Archive Control File (ACF), a data base of archived data sets used by ABR Archive Backups (DUMP TYPE=ARC) as described in [Section 51](#). If you are doing Archive Backups (especially with RECALL=YES), you will need to use FDRARCH to create an ACF and to do periodic maintenance on it.

A Control File with the same format is also used by Application Backup (DUMP TYPE=APPL) described in [Section 52](#). In some circumstances an Application Backup Control File may also need the attention of FDRARCH. In the following sections, references to the Archive Control File (ACF) will also apply to the Application Backup Control File. However, functions related to Auto-Recall and the DYNARC options will not apply.

**FDRARCH
STATEMENTS
AND
FUNCTIONS**

The FDRARCH input may contain many statements. Each function will be executed as it is encountered, except as noted below. The following FDRARCH statements provide general maintenance functions, including creation of new Archive Control Files:

FORMAT - formats a new Archive Control File.

DUMP - creates a sequential backup copy of an Archive Control File.

RESTORE - restores the backup copy of an Archive Control File created by DUMP. Can be used to move a control file to a new volume or device type, or to expand the file.

MERGE - merge multiple Archive Control Files into one.

This statement should be executed periodically to maintain your Archive Control File:

REORG - reorganizes an Archive Control File, deleting obsolete records and compressing the active records to create space for new records at the end of the file. Obsolete records can include those marked for deletion (by the DELETE statement), those that are expired (past their expiration dates), those that are flagged as having been restored by ABR, and those marked for Auto-Recall that are no longer cataloged. During the reorganization, Archive Backup files which are no longer referenced by any active record can be deleted and uncataloged.

These statements perform operations against selected records in the ARCHIVE control File:

DELETE - mark selected archive records for deletion by a subsequent REORG.

RESET - remove deletion indicators set by DELETE on selected records.

EXPIRE - mark selected records as being expired.

MODIFY - change fields within archive records in the Archive Control File (e.g., expiration dates and backup pointers).

RECATALOG - creates or updates Auto-Recall indicators for archived data sets in system catalogs, so that they can be successfully recalled.

Each of the last 5 statements (DELETE, RESET, EXPIRE, MODIFY, RECATALOG) requires that the entire Archive Control File be scanned for records to be updated. To avoid excessive passes through the file, the changes requested by these statements will be saved until another type of statement (such as REORG) or the end of the control statements is encountered; then one pass of the file will be performed, comparing each record to every request to see if an action is required. You may also explicitly request the update by entering:

UPDATE - execute all of the immediately preceding DELETE, RESET, EXPIRE, MODIFY, and RECATALOG statements.

51.50 CONTINUED . . .

**UNDOCU-
MENTED
FUNCTIONS**

FDRARCH contains options and functions which are designed for special situations or special customer requests. To avoid confusion and lengthy descriptions, some of these are not documented in this manual. For a complete list of the supported functions, execute FDRARCH with the statement:

HELP ALL(or NEWS to get a summary of recent changes)

or browse the FDRARCH member in the FDR ICL (Installation Control Library).

**OPERAND
DEFAULTS**

The defaults for many operands on FDRARCH statements may be permanently changed from the defaults shown in the descriptions which follow, by making changes in the FDR Global Option Table. [See Section 90](#) (Installation and Customization) for details; most FDRARCH options are set on ISPF panel A.I.4.6 as shown here (with its default values):

```
----- FDR INSTALLATION -- SET ABR ARCHIVE UTILITY DEFAULTS -----
COMMAND ==>

                FDRARCH UTILITY DEFAULTS
BLKF            DEFAULT NUMBER OF BLOCKS PER TRACK IN THE ARCHIVE FILE.... 2
RECS            DEFAULT NUMBER OF DATA SET ENTRIES IN THE ARCHIVE FILE.... 2000

REORG%          PERCENTAGE OF FREE SPACE IN ARCHIVE FILE TO ISSUE WARNING. 10
ARCDD           DDNAME USED BY FDRARCH TO REFERENCE ARCHIVE FILE..... ARCHIVE

DYNARC          DYNAMIC ALLOCATION OF ARCHIVE FILE BY FDRARCH ALLOWED..... NO
ERASE           ARCHIVE FILE MAY BE RE-FORMATTED (WRITTEN OVER)..... NO
ARCRESV         PROTECT ARCHIVE FILE WITH RESERVE DURING UPDATE FUNCTION.. YES

                FDRARCH ARCHIVE FILE REORGANIZATION OPTIONS
RECALL          UNCATALOG AUTO-RECALL DATA SETS THAT WERE DROPPED..... YES
UNCAT           UNCATALOG ARCHIVE TAPE FILES THAT ARE NO LONGER NEEDED.... YES
MSG             PRINT UNCATALOGED TAPE FILES THAT ARE NO LONGER NEEDED.... YES
```


51.51 FDRARCH JCL REQUIREMENTS

The following JOB control statements are required to execute the ARCHIVE maintenance utility.

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRARCH resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRARCH), region requirement (REGION=), and optional PARM= operand. Because the region requirements of FDRARCH may vary depending on the operation being performed, it is recommended that you specify REGION=0M to get the largest possible below-the-line region.

If a PARM field is specified, FDRARCH will use the data specified as the first control statement, which must be a valid FDRARCH statement. For example,

```
//FDR EXEC PGM=FDRARCH,PARM='FORMAT RECS=200000'
```

If FDRARCH is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

SYSPRINT DD STATEMENT Specifies the output message data set. This is a required DD statement and usually is a SYSOUT data set.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

ARCHIVE DD STATEMENT Specifies the Archive Control File to be processed; it must be on disk. **No DCB parameters should be specified.** The data set name must contain an index level of "ARCHIVE" in other than the first level index unless DISABLE=DSNCK is specified. The Archive Control File must reside on a disk volume that has been initialized by the ABR VTOC maintenance utility, FDRABRM, for ABR processing ([see Section 50.40](#)) unless USERINDEX=YES is specified; USERINDEX=YES is used mainly with ABR Application Backup, described in [Section 52](#).

If creating a new Archive Control File, a SPACE= parameter with secondary allocation and RLSE are supported and recommended. Estimate the number of data sets that will need to be recorded at any one time and calculate the tracks required, figuring on 292 data sets per track on a 3380 disk, and 348 per track on a 3390. However, multi-volume data sets will require one entry per volume, and ICF VSAM clusters require one entry for each cluster (on each volume) plus one entry for each component on each volume (e.g., 3 for a single volume KSDS). Be sure to add a liberal amount for obsolete entries between reorganizations. The control file **will not** take automatic secondary allocations once it has been formatted; you can increase its size by using the DUMP and RESTORE statements of FDRARCH.

For an existing Archive Control File, DISP=SHR can be used; By default, FDRARCH and FDRABR use ENQ/RESERVE macros to serialize access to the file.

```
Examples: //ARCHIVE DD DSN=FDRABR.ARCHIVE,DISP=SHR
           //ARCHIVE DD DSN=FDRABR.ARCHIVE.PROD,UNIT=3390,
           // VOL=SER=SYS001,DISP=(NEW,CATLG),
           // SPACE=(CYL,(100,10),RLSE)
```

If the ARCHIVE DD statement is not present and ENABLE=DYNARC is specified on the current statement, the Archive Control File will be dynamically allocated using the ARCDSDN data set name in the FDR global option table ([see Section 90](#)). This is normally the name of the ACF used for ABR auto-recalls. You should always use an ARCHIVE DD statement when referencing an Application Control File.

51.51 CONTINUED . . .

MERGE DD STATEMENTS If you are executing the MERGE statement of FDRARCH, you must supply 2 or more DD statements pointing to the input Archive Control Files. The DDnames are any names of your choice, and are specified in the MERGEDDNames= operand of the MERGE statement; the ARCHIVE DD statement will point to the output control file.

TAPE1 DD STATEMENT Specifies the file to be used for the backup copy of the Archive Control File. This DD statement is used to output the backup copy of the Archive Control File if the DUMP or REORG statements are executed. It is used to input a backup copy of the Archive Control File if the RESTORE statement is executed. It is not required for other functions.

NOTE: TAPE1 must point to a tape device unless the DUMPDEVICE=DISK operand is specified.

ABRWORK DD STATEMENT (Optional) used by the REORG statement if UNCAT=YES is specified or defaulted. If this DD statement is present and references a disk data set, the data set will contain a list of all the Archive Backup data sets active at the time of the REORG, their use counts, and catalog status. In addition, the DCB characteristics will be set to:

DSORG=PS, RECFM=FB, LRECL=256, BLKSIZE=4096

If this DD statement is not present, FDRARCH will use a temporary data set which will be dynamically allocated on a volume within the unit name specified by the WORKUNIT operand (default SYSALLDA).

TEMPARCH DD STATEMENT (Optional) used by the REORG statement if both SORT=YES and SIMULATE are specified. It is not required, but if present it will contain an image of what the Archive Control File would have looked like if the REORG had really been done; it can be used as input to reports in later steps. See "ARCHIVE DD STATEMENT" above for considerations on allocation.

SORT DD STATEMENTS (Optional) If you are executing the REORG statement with SORT=YES or other operands requesting sorting (see SORT= under REORG), you may need to specify DD statements required by your system SORT product (e.g. SORTLIB, SORTWKnn).

In particular, a SORTOUT DD statement will be required, with space at least as large as the Archive Control File being reorganized.

However, any or all of these SORT DD statements may be omitted if you have requested dynamic allocation of the SORT data sets via the SORTALLOC= operand of REORG.

SYSIN DD STATEMENT Specifies the control statement data set, usually an input stream or DD * data set. It can be omitted if the only control statement is specified by PARM= on the EXEC statement.

FORMAT

```
,RECS=nnnnnnnn
```

,USERINDEX=NO|YES

The FORMAT statement formats an Archive Control File for use by ABR. The output control file must be a newly-allocated, empty data set unless ENABLE=ERASE is specified.

B=

The default is 2 (half track blocking) unless overridden in the FDR Global Option Table

The options specified will be disabled or enabled for this FORMAT operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

ERASE allows an existing, formatted control file to be reformatted. If disabled, a format can only be done to a newly allocated control file.

By default, DSNCK and RESERVE are enabled, ERASE is disabled, unless overridden in the FDR Global Option Table.

Specifies the minimum number of data set entries that will be formatted on the Archive Control File, from 1 and 3000000 (3 million) .Sufficient space **must** be allocated to the Archive Control File to contain the required number of entries, otherwise FORMAT will fail.

If the SPACE= JCL parameter specifies the RLSE operand, then FDRARCH will format as many blocks as necessary to contain the required number of records and any excess space will be released.

If SPACE= does not include the RLSE operand (or SPACE= is omitted for an existing file), then FDRARCH will format all tracks of the initial allocation; if this is insufficient to hold the required number of records and a secondary allocation quantity was specified in SPACE=, FDRARCH will take secondary allocations until the space is sufficient, and will format all of the secondary tracks.

So, **FORMAT** will format all tracks allocated to the Archive Control File, with a guaranteed minimum capacity specified by **RECS=**. Once it has been formatted, it can only be expanded by **FDRARCH**.

The default is 2000 unless overridden in the FDR Global Option Table.

51.52 CONTINUED . . .

USERINDEX= YES specifies that the Archive Control File is to be formatted with an indicator that causes FDRABR to change its normal convention for naming backup files on tape or disk for any ARCHIVE (TYPE=ARC) or Application Backup (TYPE=APPL) step which references this Control File. The high-level index of such backup files will be the same as the high-level index of this Archive Control File. This is intended for use with Application Backups ([Section 52](#)) and should not be used when formatting a Control File for any other purpose.

NO specifies that the indicator will not be set and any Archive Backup files created using this Control File will follow ABR's normal naming convention (as described in [Section 51.01](#)).

Default is NO.

ISPF PANEL A.I.9 A job to execute the FDRARCH FORMAT function can be generated by the ABR ISPF install dialog panel A.I.9, shown here:

```
----- FDR INSTALLATION -- ABR ARCHIVE FILE INITIALIZATION -----
COMMAND ==>

SUBMIT - SUBMIT INITIALIZATION JOB      CANCEL - EXIT IMMEDIATELY

ARCHIVE DATA SET NAME   ==> FDRABR.ARCHIVE
STANDARD DATA SET NAME  ==> YES
MAX DATA SET ENTRIES    ==> 10000
BLOCKS PER TRACK         ==> 2

VOLUME SERIAL NUMBER     ==>

FDR PROGRAM LIBRARY NAME ==> 'IDP.MODFDR53'
VOLUME SERIAL            ==>

SYSPRINT SYSOUT CLASS    ==> *

JOB STATEMENT INFORMATION:
==> //useridA JOB (ACCOUNT),'NAME',
==> //          NOTIFY=userid
==> //*
==> /**
```

This will allocate, catalog and format the Archive Control File shown at the top of the screen:

Archive Data Set Name - name of the file to be created, quotes are not required.

Standard Data Set Name - YES corresponds to ENABLE=DSNCK

Max Data Set Entries - corresponds to REC=

Blocks per Track - corresponds to B=

Volume Serial Number - specify the location, which must be an ABR-initialized volume.

The size of the file will be automatically calculated based on the values you provided.

The FDR program library name will probably be already filled in. If necessary update it and provide the volume serial if it is not cataloged. If necessary, update the JOB statement shown at the bottom to meet your installation's standards.

When done, enter SUBMIT on the command line to submit the job.

51.53 FDRARCH DUMP STATEMENT

DUMP**DUMPDEVICE=DISK|TAPE****,DISABLE=(option,option,...)****,ENABLE=(option,option,...)****DUMP
STATEMENT**

The DUMP statement causes a formatted backup copy of the Archive Control File to be created. This backup is not in FDR format, and can only be restored by the RESTORE statement of FDRARCH. The Archive Control File may be dumped by FDRABR in FDR format at the end of each dump TYPE=ARC operation.

The DUMP and RESTORE statements of FDRARCH can be used to move the Archive Control File to a different device type (such as 3380 to 3390) or to change its size or blocking factor. FDRCOPY or FDRDSF DUMP/RESTORE can also be used to move the Archive Control File, but will not expand or reblock the file.

OPERANDS**DUMPDEVICE=**

Specifies the device type that is to be used as a DUMP medium.

DISK – permits the DUMP data set to be placed on either disk or tape devices.

TAPE – restricts the DUMP data set to tape devices only.

The default is TAPE.

**DISABLE=
ENABLE=**

The options specified will be disabled or enabled for this DUMP operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

DSNCK validates the data set name of the Archive Control File. It requires that some index level in the name, not the first, is ARCHIVE.

DYNARC dynamically allocates the common Archive Control File named in the ARCDN option of the FDR Global Option Table if the ARCHIVE DD statement is not present.

RESERVE protects the Control File with a RESERVE macro against the volume on which it resides.

By default, DSNCK and RESERVE are enabled, DYNARC is disabled, unless overridden in the FDR Global Option Table.

51.54 FDRARCH RESTORE STATEMENT

RESTORE

B=n,DISABLE=(option,option,...)
,ENABLE=(option,option,...)
,DUMPDEVICE=DISK|TAPE,RECS=nnnnnnnn

**RESTORE
STATEMENT**

The RESTORE statement creates an Archive Control File from a backup copy created by the FDRARCH DUMP statement. If the data set to which FDRARCH is restoring is newly allocated, it does not have to be pre-initialized with the FORMAT statement because the RESTORE statement includes the functions of FORMAT.

The DUMP and RESTORE statements of FDRARCH can be used to move the Archive Control File to a different device type (such as 3380 to 3390) or to change its size or blocking factor. The RESTORE statement of FDRARCH cannot be used to restore a backup of the Archive Control File that was created by FDRABR or FDRDSF; that backup can only be restored by FDRABR or FDRDSF.

OPERANDS**B=**

Specifies the number of blocks to be placed on each track, from 2 to 5 . FDRARCH will automatically calculate the best blocksize for this device type.

The default is 2 (half-track blocking) unless overridden in the FDR Global Option Table.

DUMPDEVICE=

Specifies the device type that is eligible to be used as a RESTORE medium.

DISK— accepts RESTORE data from either tape or disk devices.

TAPE— restricts the RESTORE data to tape devices only.

The default is TAPE.

**DISABLE=
ENABLE=**

The options specified will be disabled or enabled for this RESTORE operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

DSNCK validates the data set name of the Archive Control File. It requires that some index level in the name, not the first, is ARCHIVE.

DYNARC dynamically allocates the common Archive Control File named in the ARCDN option of the FDR Global Option Table if the ARCHIVE DD statement is not present.

ERASE allows an existing, formatted control file to be restored. If disabled, a restore can only be done to a newly allocated control file.

RESERVE protects the Control File with a RESERVE macro against the volume on which it resides.

By default, DSNCK and RESERVE are enabled, ERASE and DYNARC are disabled, unless overridden in the FDR Global Option Table.

RECS=

Specifies the minimum number of data set entries that will be formatted during the restore of the Archive Control File, from 1 and 3000000 (3 million) . See RECS= in [Section 51.52](#) for details.

The default is 2000 unless overridden in the FDR Global Option Table.

51.55 FDRARCH MERGE STATEMENT

MERGE **MERGEDDNames=(ddname,ddname,...)** **,DISABLE=(option,option,...)**
 ,B=n **,ENABLE=(option,option,...)**
 ,RECATALOG=YES|NO
 ,CATERROR=IGNORE|PURGE|RETAIN
 ,RECS=nnnnnnnn

MERGE STATEMENT The MERGE statement merges two or more Archive Control Files into a single new Archive Control File. Up to 16 Archive Control Files may be merged in a single execution of the MERGE statement. All records regardless of their status (DELETED, EXPIRED, RESTORED, etc.) will be copied. If records must be removed prior to merging the files, use the DELETE and REORG statements before merging the files. Records in the merged control file will be in the order of their creation date (date of ARCHIVE); for a given date records will be merged first from the first ddname specified, then from the second, etc.

Catalog entries for Auto-Recall data sets contain a pointer to the block in the ACF where the search for the data set is to start. Since ABR searches the ACF backwards, this will still work even if reorganization moves data set records to lower block number. But If records in the input Archive Control Files were archived for Auto-Recall, a data set record in the merged files may end up in a **higher** block number than that in the catalog entry, causing Auto-Recall to fail. By default, MERGE will update the catalog entry when this is detected. However, this is the correct thing to do **only** if the merged ACF will become the common Archive Control File for Auto-Recall on this system. If it is **not** to become the common file, specify RECATALOG=NO.

MERGE can be used to merge the Archive Control Files that were previously used by separate ABR systems, but which are now to be run as a single system. Since only one control file can be used for Auto-Recall, this merge may be required when combining data centers, such as after a company acquisition. However, it is generally possible to merge control files **only** if the same ABR high-level index (e.g., FDRABR) was used on both systems. **Innovation strongly recommends that you call for guidance before attempting such a MERGE.**

OPERANDS	MERGEDDNames=	Specifies the DD names (from 2 to 16) to be read as input when merging Archive Control Files into a single Archive Control File. MERGEDDNames= is required.
	B=	Specifies the number of blocks to be placed on each track, from 2 to 5 . FDRARCH will automatically calculate the best blocksize for this device type. The default is 2 (half-track blocking) unless overridden in the FDR Global Option Table.
	CATERROR=	Specifies the processing required if an error is encountered while accessing an ICF catalog. IGNORE ignores the error and treats it like a "not-found" condition (i.e., data set is not cataloged) unless the result would be creation of a catalog entry in which case it is treated like CATERROR=RETAIN. PURGE causes the entry to be purged from the Archive Control File, if that is an available option. If not, it is treated like CATERROR=IGNORE. RETAIN stops processing on the entry and retains it for reprocessing after the error within the catalog has been corrected. The default is RETAIN.

51.55 CONTINUED . . .

DISABLE=
ENABLE=

The options specified will be disabled or enabled for this MERGE operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

DSNCK validates the data set name of the Archive Control File. It requires that some index level in the name, not the first, is ARCHIVE.

DYNARC dynamically allocates for output the common Archive Control File named in the ARCDN option of the FDR Global Option Table if the ARCHIVE DD statement is not present.

ERASE allows an existing, formatted control file to be used as output for the MERGE. If disabled, a MERGE can only be done to a newly allocated control file.

RESERVE protects the Control File with a RESERVE macro against the volume on which it resides.

By default, DSNCK and RESERVE are enabled, ERASE and DYNARC are disabled, unless overridden in the FDR Global Option Table.

RECATALOG=

YES – Specifies that for any entries in any input Archive Control File which are flagged for Auto-Recall and are still cataloged for Auto-Recall, MERGE will update the pointers in the catalog entries for those data sets to point to the proper location in the merged control file.

RECATALOG=YES should be used when the merged control file is to become the common Archive Control File for Auto-Recall; otherwise Auto-Recalls may fail.

NO – specifies that Auto-Recall catalog entries are not to be corrected.

RECATALOG=NO should be used when the merged control file is not to become the common Archive Control File (such as a test of the MERGE function). If necessary, you can use the RECATALOG statement to correct the Auto-Recall entries if the merged file later becomes the common file.

The default is YES.

RECS=

Specifies the minimum number of data set entries that will be formatted in the output Archive Control File, from 1 and 3000000 (3 million) . See RECS= in [Section 51.52](#) for details.

Sufficient space **must** be allocated to the Archive Control File to contain the required number of entries, otherwise MERGE will fail.

The default is the number of records resulting from the MERGE operation.

51.56 FDRARCH REORG STATEMENT

REORG	ADAYS=nnnnn	,RECALL= <u>YES</u> NO
	,CATERROR=IGNORE PURGE <u>RETAIN</u>	,RESTORE=IGNORE PURGE RETAIN
	,CNTRESV=nn	,SEARCH= <u>CVA</u> F OBTAIN NO
	,DELETE=IGNORE PURGE RETAIN	,SIMULATE
	,DISABLE=(option,option,..) ,ENABLE=(option,option,..)	,SORT=YES <u>NO</u>
	,DUMPDEVICE=DISK <u>TAPE</u>	,SORTALLOC= <u>NO</u> SORTLIB SORTMSG SORTOUT SORTWORK YES
	,EXPIRE=IGNORE PURGE RETAIN	,SORTCORE=nnnnnnnn
	,LOG NOLOG	,SORTLIB=dsname
	,MAXGENERATIONS=nnnnn	,SORTMSGDDNAME=ddname
	,MAXOCCURRENCES=nnnnn	,SORTPFX=cccc <u>SORT</u>
	,MAXUNCATALOG=nnnnn	,UNCAT=NO <u>YES</u>
	,MSG=NO <u>YES</u>	,WORKDDNAMES=n
		,WORKUNIT=unitname <u>SYSALLDA</u>

REORG STATEMENT The REORG statement first creates a backup copy of the Archive Control File on TAPE1, similar to the DUMP statement. Then it restores the backup back to the ACF, but while doing so it removes entries that are no longer required. You can specify various criteria for determining the unneeded entries. The space occupied by the removed records is recovered and there will be additional unused blocks at the end of the file for the addition of new entries.

During this reorganization, REORG may perform other support functions:

- All Archive Backup files no longer referenced within the Archive Control File may optionally be deleted (if on disk) and uncataloged from the ABR catalog.
- Any data sets within the Archive Control File that are marked for auto recall may also be optionally uncataloged if they are removed.
- Any data sets marked for auto recall which are actually on disk (no longer archived) can be recataloged normally.

51.56 CONTINUED . . .

**REORG
STATEMENT
(continued)**

REORG can be a high-overhead operation because of the large number of catalog accesses (SVC 26) and VTOC accesses required for data sets flagged for auto recall. Improvements of the performance of REORG on large control files have been made in recent releases of ABR. REORG can optionally use a SORT to improve performance on large control files (typically over 100,000 records), sorting the ARCHIVE records and then accessing catalogs sequentially. You can request the SORT-based REORG by specifying SORT=YES or it will be automatically invoked if certain options are specified (such as MAXOCCURRENCES=). However, if SORT is invoked, you may need to add JCL statements required by your installation's SORT product or add the SORTALLOCATE= operand to REORG to request that FDRARCH dynamically allocate SORT files for you. Some SORT products may require neither.

NOTE: a FDRARCH REORG of your Archive Control File (the common one used for Auto-Recalls) on a regular basis is a necessity. The frequency of the REORG depends on your ARCHIVE activity and the size of your control file; most installations do the REORG weekly.

WARNING: *once a REORG has begun, it should be allowed to run to completion if at all possible. Cancellation of a REORG usually causes no problems but the potential exists for damage. Because of this, FDRARCH now includes CANCEL protection during REORG: any attempt to cancel the REORG job will be intercepted and the operator will be offered the option of ignoring the CANCEL and continuing the REORG, allowing the CANCEL to take place, or restoring the Archive Control File to its original state; this last option will not undo any catalog updates that have already taken place.*

OPERANDS**ADAYS=**

Specifies that data sets archived with the last "n" days (from 0 to 32000) are to be protected from REORG processing. They will not be purged regardless of other REORG operands.

The default is 0 (no data sets protected).

CATERROR=

Specifies the processing required if an error is encountered while accessing an ICF catalog.

IGNORE ignores the error and treats it like a "not-found" condition (i.e., data set is not cataloged) unless the result would be creation of a catalog entry, in which case it is treated like CATERROR=RETAIN.

PURGE causes the entry to be purged from the Archive Control File, if that is an available option. If not, it is treated like CATERROR=IGNORE.

RETAIN stops processing on the entry and retains it for reprocessing after the error within the catalog has been corrected.

The default is RETAIN.

CNTRESV=

Specifies the number of Archive Control File blocks (from 1 to 99) to be processed between times when the RESERVE on the volume containing the control file is released. This allows other processors to access the volume at intervals during the reorganization.

The default is 10.

51.56 CONTINUED . . .

DELETE=

Specifies the processing required for data sets previously marked by the DELETE statement of FDRARCH for removal from the Archive Control File:

IGNORE do not test the DELETE indicator

PURGE drop the data set from the Archive Control File unless prevented by other selection criteria (e.g., EXPIRE=RETAIN)

RETAIN retain the data set even if other selection criteria calls for its removal

The default is PURGE but if EXPIRE= and/or RESTORE= is also specified, the default is IGNORE.

51.56 CONTINUED . . .

DISABLE=
ENABLE=

The options specified will be disabled or enabled for this REORG operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

BYPASSCATERR does not report catalog processing errors detected by FDRARCH.

DSNCK validates the data set name of the Archive Control File. It requires that some index level in the name, not the first, is ARCHIVE.

DYNARC dynamically allocates the common Archive Control File named in the ARCDN option of the FDR Global Option Table if the ARCHIVE DD statement is not present.

EXPANDREASON prints expanded drop reasons during ARCHIVE reorganize.

IFNOTAUTOREC deletes **all** entries that are not marked for auto recall, regardless of other operands.

IFNOTCAT deletes **all** entries that have been archived for auto recall but are not currently cataloged for auto recall, regardless of other operands. If a data set was archived, recalled, and re-archived, the record of the older copy will also be deleted by IFNOTCAT. The PRINTxxxx options below are effective only if IFNOTCAT is enabled. IFNOTCAT allows users to delete data sets from the Archive Control File simply by uncataloging them. It also deletes GDG generations which were archived but are no longer cataloged.

PRINTALIAS prints all ALIAS catalog entries processed by FDRARCH.

PRINTALL prints all catalog entries processed by FDRARCH, even those not selected.

PRINTSELECT prints all catalog entries selected for processing by FDRARCH.

RESERVE protects the Control File with a RESERVE macro against the volume on which it resides.

SMSEXPIRE requests that a COPY 1 backup on disk will be deleted when the COPY 1 expiration date has been reached on all entries which point to that backup. If disabled, the COPY 1 on disk is deleted only when all entries which point to it are being deleted from the control file (usually because **both** copies have reached their expirations). This is used to support the SMSEXPIRE= option of ABR (See Section 51.05 and 70).

TMS indicates that an automated tape management system, such as CA-1, CA-TLMS, or DFSMSrmm is in use in the system. Enabling this option means FDRARCH will only uncatalog and/or recatalog an Archive Backup file if the expiration date recorded in the control file is 99000 ("catalog control" in most tape management systems).

By default, DSNCK and RESERVE are enabled and all others are disabled, unless overridden in the FDR Global Option Table.

51.56 CONTINUED . . .

DUMPDEVICE=	<p>Specifies the device type that is eligible to be used as TAPE1 (dump output).</p> <p>DISK Permits the DUMP data set to be placed on either disk or tape devices.</p> <p>TAPE Restricts the DUMP data set to tape devices only.</p> <p>The default is TAPE.</p>
EXPIRE=	<p>Specifies the processing required for any expired data sets, entries whose expiration date has been passed. If the data set has 2 backup copies, the expiration for both COPY 1 and COPY 2 must be passed. FDRARCH will not consider data sets with expirations of 99.000, 99.365, or 99.366 to be expired.</p> <p>IGNORE do not test the expiration dates.</p> <p>PURGE drops the data set from the Archive Control File unless prevented by other selection criteria (e.g., RESTORE=RETAIN)</p> <p>RETAIN retains the data set even if other selection criteria calls for its removal</p> <p>The default is PURGE but if DELETE= or RESTORE= is specified the default is IGNORE.</p>
LOG NOLOG	<p>LOG specifies that the messages detailing the entries dropped from the Archive Control File are to be printed.</p> <p>NOLOG suppresses detail messages and prints only a summary.</p> <p>The default is LOG.</p>
MAXGENERATIONS=	<p>Specifies the maximum number of archived generations of any given generation data group (GDG) that will reside in the Archive Control File at the end of the reorganization. The number may be any number from 1 to 32000.</p> <p>The default is the GDG status is not tested.</p>
MAXOCCURRENCES=	<p>Specifies the maximum number of archived copies of any given data set that will reside in the Archive Control File at the end of the reorganization. The number may be any number from 1 to 32000.</p> <p>The default is the maximum occurrences of a given data set is not tested.</p>
MAXUNCATALOG=	<p>If RECALL=YES is specified or defaulted, and a SORT-based REORG is being done (SORT=YES or other options implying SORT), this specifies the maximum number (from 1 to 32000) of auto recall data sets which will uncataloged in one SORT pass. If more than this number of uncatalogs is required, additional SORT passes will be required.</p> <p>The default is 8000.</p>

51.56 CONTINUED . . .

MSG=	<p>NO – Do not print uncatalog processing messages.</p> <p>YES – print messages for every Archive Backup file uncataloged during this REORG.</p> <p>The default is YES unless overridden in the FDR Global Option Table.</p>
RECALL=	<p>Specifies the processing requirements for those Archive Control File records marked for auto recall.</p> <p>NO – Ignore the auto recall indicator.</p> <p>YES – Uncatalog the special auto recall entry in the system catalog whenever a data set flagged for auto recall is removed from the Archive Control File during REORG.</p> <p>The default is YES.</p>
RESTORE=	<p>Specifies the processing required for any restored data sets (data sets that are marked as having been successfully restored by FDRABR) that are encountered.</p> <p>IGNORE do not test the RESTORED indicator</p> <p>PURGE drops the data set from the Archive Control File unless prevented by other selection criteria (e.g., EXPIRE=RETAIN)</p> <p>RETAIN retains the data set even if other selection criteria calls for its removal</p> <p>The default is IGNORE.</p>
SEARCH=	<p>Specifies if and how REORG is to verify that data sets flagged for auto recall are accurately cataloged for recall (not on disk).</p> <p>CVAF uses Common VTOC Access Facility macros.</p> <p>OBTAIN uses the OBTAIN DADSM macro</p> <p>NO bypasses this verification.</p> <p>If CVAF or OBTAIN is specified, for every Auto-Recall entry which is being removed from the control file, if the original DASD volume recorded in the control file is online and of the right type (e.g., 3390), REORG will access the VTOC to see if the data set is on that volume. If it is, it will be recataloged to that volume instead of removing the catalog entry.</p> <p>The default is CVAF.</p>
SIMULATE	<p>Specifies that the reorganization of the Archive Control File is only to be simulated. No entries are uncataloged. The Archive Control File is not rewritten. If the TEMPARCH DD statement (see Section 51.51) is present, it will contain a copy of what the control file would have looked like after the reorganization.</p> <p>The default is to actually reorganize the Archive Control File.</p> <p>Warning: Innovation strongly recommends that any new or updated REORG jobstream be run in simulation mode first to insure that it will delete the intended entries.</p>

51.56 CONTINUED . . .

SORT=

YES – the SORT-based version of the REORG statement is to be used

NO – the SORT-based version of the REORG statement is not to be used unless one or more operands are specified that are requiring sorting; these include:

MAXGENERATIONS, MAXOCCURRENCES, SORTALLOC, SORTCORE, SORTLIB, SORTMSG, SORTMSGDDNAME, SORTPFX, WORKDDNAMES

The default is NO.

For best performance, Innovation strongly recommends SORT=YES for most REORGs.

SORTALLOC=

Specifies if FDRARCH is to dynamically allocate sort related DD statements.

NO – do not dynamically allocate SORT related data sets. If sorting is required, any necessary DD statements must be included in the JCL or dynamically allocated by your SORT product.

SORTLIB – dynamically allocate the SORTLIB data set using the value in the operand SORTLIB= for the dsname. Users of the SYNC SORT product should see the note under the SORTLIB= operand.

SORTMSG – dynamically allocate the SORT message output to SYSOUT using the value in the operand SORTMSGDDNAME= for the ddname.

SORTOUT – dynamically allocate a temporary SORTOUT data set using the value in the operand WORKUNIT= as the unitname.

SORTWORK – dynamically allocate the number of SORTWKnn data sets specified in the operand WORKDDNAMES, using the value in the operand WORKUNIT= as the unitname.

YES – dynamically allocate all of the above.

Multiple values may be specified for this operand in parenthesis separated by commas. For example,

```
REORG SORTALLOC=(SORTWORK, SORTOUT)
```

specifies that SORT work and output DD statements are to be dynamically allocated to temporary data sets.

The default is NO.

SORTCORE=

Specifies the amount of storage (from 10000 to 8000000, in bytes) SORT is to use.

The default is taken from the FDR Global Option Table and is usually 100000. It can be changed on ISPF panel A.I.4.9 ([see Section 90](#)).

51.56 CONTINUED . . .

SORTLIB=	<p>Specifies the data set name to be allocated to the DDNAME SORTLIB.</p> <p>The default is SYS1.SORTLIB.</p> <p>NOTE for SYNCSORT users: The SYNCSORT SORT product does not require a SORTLIB, so SYS1.SORTLIB may not exist on your system. If SORTALLOC=ALL is specified, FDRARCH will attempt to allocate it, and it may fail. To circumvent this, you can either create an empty PDS called SYS1.SORTLIB, or override the SORTLIB= operand to specify some other PDS; this dummy data set will not actually be used.</p>
SORTMSGDDNAME=	<p>Specifies the DDNAME to be used by the program SORT if messages are to be printed.</p> <p>The default is SYSOUT.</p>
SORTPFX=	<p>Specifies the DDNAME prefix to be used by the program SORT if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.</p> <p>The default is taken from the FDR Global Option Table and is usually "SORT".</p>
UNCAT=	<p>Specifies whether Archive disk or tape backup files which are no longer referenced in the Archive Control File due to this REORG (all data sets which were archived into those files have been PURGED) are to be scratched and/or uncataloged. If the TMS option is in effect (see ENABLE=), tapes will be uncataloged only if they have an expiration date of 99.000 (catalog control).</p> <p>NO – Do not uncatalog Archive Backup files that are no longer referenced.</p> <p>YES – Uncatalog Archive Backup files that are no longer referenced. The default is YES.</p>
WORKDDNAMES=	<p>Specifies the number of SORT work DDnames to allocate, from 1 to 5.</p> <p>The default is 3.</p>
WORKUNIT=	<p>Specifies the unit name (1 to 8 characters) to use when dynamically allocating any of the temporary work files, including sort work files if requested by SORTALLOC=. It must be a value valid for UNIT= in JCL, and the volumes included on those units must include some in STORAGE or PUBLIC status for the allocation to be successful.</p> <p>The default is SYSALLDA which is valid on all MVS systems and includes all DASD devices.</p>

51.57 FDRARCH DELETE STATEMENT

```

DELETE      ADATE=yyddd|yyyyddd/hhmmss      ,REAL|SIMULATE
            ADAYS=nnnnn                      ,TAPEFILE=nnnn
            ,DSN=(dsname,...,dsname)          ,TAPEVOL=(vvvvvv,...,vvvvvv)
            ,DSG=(dsgroup,...,dsgroup)        ,VOL=(vvvvvv,...,vvvvvv)
            ,XDSN=(mask,...,mask)
            ,EXPDATE=yyddd|yyyyddd
            ,EXPDAYS=nnnnn

```

**DELETE
STATEMENT**

The DELETE statement is used to mark records within the Archive Control File for deletion. The statement may appear more than once in the control statement input if you need to delete records based on several sets of criteria. The actual deletion will not take place until the next reorganization (REORG) of the control file; up to that time you may use the RESET statement to turn off the DELETE indicator.

The DELETE parameters are saved for later processing. Please review [Section 51.50](#) to understand when the Archive Control File will be scanned and the DELETE indicator set in the selected records.

Many of the operands of DELETE will accept comparison operators in addition to a simple equal (=). The operators are:

= or .EQ.	Equal
≠ or .NE.	Not Equal
< or .LT.	Less Than
> or .GT.	Greater Than
<= or .LE.	Less Than or Equal
>= or .GE.	Greater Than or Equal

In most cases, you may also specify an operand more than once, such as:

```
ADAYS>5,ADAYS<20
```

to select a range of values.

A record must meet **all** of the selection criteria specified to be marked for deletion.

OPERANDS**ADATE=**

Specifies the date of the ARCHIVE run of the records to be selected. Only data sets archived on this date will be marked. See the note on operators above to specify comparisons and ranges.

The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).

For data sets that were archived with RECALL=NO (no Auto-Recall), including Application Backups, the control file contains a time of day as well as a date. In this case, you can select by date and time, e.g., ADATE=97123/140522 or ADATE=1997123/140522.

ADAYS=

Specifies that records that were archived "nnnnn" days ago (from 1 to 32000) are to be selected. See the note on operators above to specify comparisons and ranges.

51.57 CONTINUED . . .

DSN=	Specifies one or more data set names (up to 44 characters in length) to be selected. Multiple names may be specified in parentheses, separated by commas. See the note on operators above to specify comparisons and ranges.
DSG=	Specifies one or more data set name prefixes (up to 44 characters in length) to be selected. Multiple prefixes may be specified in parentheses, separated by commas. Only Archive Control File entries whose name begins with one of the prefixes will be marked. See the note on operators above to specify comparisons and ranges. The XDSN= operand, described below, offers a much more flexible way of selecting data sets. Note: DSN= and DSG= may be repeated and/or intermingled.
EXPDATE=	Specifies the expiration date of the ARCHIVE records to be selected. If both a COPY 1 and COPY 2 backup exist for the data set, the higher of the two expirations is tested. See the note on operators above to specify comparisons and ranges. The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).
EXPDAYS=	Specifies that records whose expiration date is within "nnnnn" days of today's date are to be selected. The value can be from 1 to 32000 and can be positive (e.g., EXPDAYS=5 or EXPDAYS=+5 means 5 days from today) or negative (e.g., EXPDAYS=-5 means 5 days ago). It will never consider data sets whose expiration date is 99.000, 99.365 or 99.366 to be expired. See the note on operators above to specify comparisons and ranges.
REAL SIMULATE	SIMULATE causes the statement to operate in simulation mode, so that you can verify that the parameters you have specified will have the correct results. It will report on what records it would modify if it was not in simulation mode, but will not update the control file. REAL (the default) makes the desired modifications in the selected records.
TAPEFILE=	Specifies a file sequence number (1 to 4095) of the archive tape data set that contains the entry that is to be selected. Ignored unless TAPEVOL= is also specified.

51.57 CONTINUED . . .

TAPEVOL= Specifies a volume serial number of the archive tape data set that contains the entry that is to be selected. The volume serial number will be matched against both the copy 1 and copy 2 volume serial number lists.

Multiple serials can be specified by enclosing the list in parentheses and separate them by commas. However, the complete list of serials will be compared to every record in the Archive Control File. If a given ARCHIVE run used multiple output tapes, you may have some data sets in a backup file entirely contained on one tape, some entirely on the second tape, and some that are in a file that crossed between the two tapes. You may need multiple DELETE statements with, for example:

```
TAPEVOL=000001
TAPEVOL=(000001,000002)
TAPEVOL=000002
```

to select all the entries on those tapes.

VOL= Specifies the volume serial number of the disk from which the data sets to be selected were archived. Only those records having a matching originating volume serial number will be selected. Multiple serials may be specified by enclosing them in parentheses separated by commas. Volume serial prefixes can be specified by following the prefix with an asterisk, e.g., VOL=TSO*. See the note on operators above to specify comparisons and ranges.

XDSN= Specifies one or more data set name masks, as defined in [Section 80.14](#). The mask allows for very flexible specification for the data set names to be selected. Multiple masks may be specified in parentheses, separated by commas. Only Archive Control File entries whose name matches one of the masks will be selected. See the note on operators above; XDSN accepts only equal (=) or not equal (≠) comparisons.

Examples: XDSN=TSO.**.LIST - all data sets whose first index is TSO and whose last is LIST

XDSN=ABC+++.** - all data sets whose first index is ABR plus 3 numeric digits

51.58 FDRARCH RESET STATEMENT

RESET	ADATE =yyddd yyyyddd/hhmmss ADAYS =nnnnn ,DSN =(dsname,...,dsname) ,DSG =(dsgroup,...,dsgroup) ,XDSN =(mask,...,mask) ,EXPDATE =yyddd yyyyddd ,EXPDAYS =nnnnn	,<u>REAL</u> SIMULATE ,TAPEFILE =nnnn ,TAPEVOL =(vvvvvv,...,vvvvvv) ,VOL =(vvvvvv,...,vvvvvv)
--------------	--	--

RESET STATEMENT The RESET statement is used to reset the deletion indicator in records within the Archive Control File. The statement may appear more than once in the control statement input if you need to reset records based on several sets of criteria.

The RESET parameters are saved for later processing. Please review [Section 51.50](#) to understand when the Archive Control File will be scanned and the DELETE indicator reset in the selected records.

OPERANDS The operands of the RESET statement are identical to those of the DELETE statement and are documented in [Section 51.57](#).

51.59 FDRARCH EXPIRE STATEMENT

EXPIRE	ADATE= yyddd yyyyddd/hhmmss ADAYS= nnnnn ,DSN= (dsname,...,dsname) ,DSG= (dsgroup,...,dsgroup) ,XDSN= (mask,...,mask) ,EXPDATE= yyddd yyyyddd ,EXPDAYS= nnnnn	,REAL SIMULATE ,TAPEFILE= nnnn ,TAPEVOL= (vvvvvv,...,vvvvvv) ,VOL= (vvvvvv,...,vvvvvv)
---------------	--	---

EXPIRE STATEMENT The EXPIRE statement is used to mark records within the Archive Control File as expired by setting their expiration date (in both copies, if required) to today's date.. The statement may appear more than once in the control statement input if you need to expire records based on several sets of criteria.

The EXPIRE parameters are saved for later processing. Please review [Section 51.50](#) to understand when the Archive Control File will be scanned and the expiration date set in the selected records.

OPERANDS The operands of the EXPIRE statement are identical to those of the DELETE statement and are documented in [Section 51.57](#).

51.60 FDRARCH MODIFY STATEMENT

MODIFY	REAL <u>SIMULATE</u>	These operand specify the modifications to be made to the selected records; specify one or more of them:
	,LOG <u>NOLOG</u>	
	,DETAIL <u>NODETAIL</u>	,NEWDEV=devtype
	These operands control selection of records to be modified; specify one or more of them:	,NEWEXDATE=yyddd yyyyddd ,NEWEXPIRE=ADATE EXPIRE TODAY ,DAYS=-nnnn ,DAYS+=nnnn
	,ADATE=yyddd yyyyddd/hhmmss ,ADAYS=nnnnn	,NEWTAPEVOL=(vvvvvv,...,vvvvvv) ,NEWTAPEFILE=nnnn
	,DSN=(dsname,...,dsname) ,DSG=(dsgroup,...,dsgroup) ,XDSN=(mask,...,mask)	,RESET=(option,option,..)
	,EXPDATE=yyddd yyyyddd ,EXPDAYS=nnnnn	,SET=(option,option,..)
	,TAPEVOL=(vvvvvv,...,vvvvvv)	,TAPECOPY=1 2 BOTH
	,TAPEFILE=nnnn	
	,VOL=(vvvvvv,...,vvvvvv)	

MODIFY STATEMENT The MODIFY statement is used to make changes to selected data set records within the Archive Control File. Entries to be modified may be selected using any combination of the operands in the first column above. The statement may appear more than once in the control statement input if you need to modify records based on several sets of criteria. In those selected entries, the expiration date and/or the archive backup volume serials, file number, device type and certain flags may be modified.

The MODIFY statement is commonly used to extend the expiration of archived data sets, by changing the expiration date recorded in the Archive Control File.

NOTE: It is the user's responsibility to make any required corresponding changes in the expiration dates of the Archive Backup tapes recorded by a tape management system.

MODIFY may also be used to change the volume serial(s), file number, and device type of an Archive Backup. This may be used after an FDRTSEL or FDRTCOPY run ([See Section 60](#)) if it failed to update the Archive Control File, to avoid having to recopy the tape.

51.60 CONTINUED . . .

**MODIFY
STATEMENT
(continued)**

The MODIFY parameters are saved for later processing. Please review [Section 51.50](#) to understand when the Archive Control File will be scanned and the modifications made in the selected records.

Many of the operands of MODIFY will accept comparison operators in addition to a simple equal (=). The operators are:

=	or .EQ.	Equal
≠	or .NE.	Not Equal
<	or .LT.	Less Than
>	or .GT.	Greater Than
<=	or .LE.	Less Than or Equal
>=	or .GE.	Greater Than or Equal

In most cases, you may also specify an operand more than once, such as:

ADAYS>5,ADAYS<20

to select a range of values.

A record must meet **all** of the selection criteria specified to be selected for modification.

WARNING: MODIFY should always be run in SIMULATION mode (without the "REAL" operand) first to be sure that the expected changes will be made. We also recommend that a backup of the Archive Control File be taken first so that changes can be easily backed out. Please contact INNOVATION if you have questions on the use of the MODIFY statement.

OPERANDS**ADATE=**

Specifies the date of the ARCHIVE run of the records to be selected. Only data sets archived on this date will be marked. See the note on operators above to specify comparisons and ranges.

The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).

For data sets that were archived with RECALL=NO (no Auto-Recall), including Application Backups, the control file contains a time of day as well as a date. In this case, you can select by date and time, e.g., ADATE=97123/140522 or ADATE=1997123/140522.

ADAYS=

Specifies that data sets that were archived "nnnnn" days ago (from 1 to 32000) are to be selected. See the note on operators above to specify comparisons and ranges.

**DETAIL
NODETAIL**

DETAIL prints messages showing every field that was updated in each selected data set. DETAIL implies the LOG operand.

NODETAIL suppresses the detail messages.

The default is NODETAIL.

DSN=

Specifies one or more data set names (up to 44 characters in length) to be selected. Multiple names may be specified in parentheses, separated by commas. See the note on operators above to specify comparisons and ranges.

51.60 CONTINUED . . .

DSG=	<p>Specifies one or more data set name prefixes (up to 44 characters in length) to be selected. Multiple prefixes may be specified in parentheses, separated by commas. Only Archive Control File entries whose name begins with one of the prefixes will be marked. See the note on operators above to specify comparisons and ranges. The XDSN= operand, described below, offers a much more flexible way of selecting data sets.</p> <p>Note: They may be repeated and/or intermingled.</p>
EXPDATE=	<p>Specifies the expiration date of the ARCHIVE records to be selected. If both a COPY 1 and COPY 2 backup exist for the data set, the higher of the two expirations is tested. See the note on operators above to specify comparisons and ranges.</p> <p>The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).</p>
EXPDAYS=	<p>Specifies that records whose expiration date is within "nnnnn" days of today's date are to be selected. The value can be from 1 to 32000 and can be positive (e.g., EXPDAYS=5 or EXPDAYS=+5 means 5 days from today) or negative (e.g., EXPDAYS=-5 means 5 days ago). It will never consider data sets whose expiration date is 99.000, 99.365 or 99.366 to be expired. See the note on operators above to specify comparisons and ranges.</p>
LOG NOLOG	<p>LOG specifies that messages listing the data sets that were updated are to be printed. If the DETAIL operand is also specified, the individual fields updated in each record are also listed.</p> <p>NOLOG suppresses data set and detail messages.</p> <p>The default is NOLOG unless overridden on a UPDATE statement (see Section 51.62).</p>
NEWDEV=	<p>Specifies that the selected entries are to be modified to indicate a new device type for the archive backup file. NEWDEV= will be ignored unless NEWTAPEVOL= is also specified. Valid values are:</p> <p>Disk: 3380, 3390</p> <p>Tape: 3400-6 (6250 BPI round tape), 3480, 3480X (with IDRC), 3490 (for 3490E), 3590-1 (Magstar)</p>

51.60 CONTINUED . . .

NEWEXDATE=
NEWEXPIRE=
DAYS=-
DAYS+=

Specifies that the selected entries are to be modified with a new expiration date for the archive backup files. The new date can be an explicitly specified Julian date (NEWEXDATE=yyddd or yyyyddd) or it can be calculated (NEWEXPIRE= and either DAYS=- or DAYS+=). NEWEXPIRE specifies the base date for the calculation and can have values:

ADATE – the date that the selected data set was originally archived.

EXPIRE – the current expiration date of the data set.

TODAY – the current date.

DAYS=- or DAYS+= should be specified with NEWEXPIRE= and specify the number of days to be subtracted from or added to the base date to arrive at the new expiration. If NEWEXPIRE= is specified without DAYS+= or DAYS=-, the default is DAYS+=30 (add 30 days to base date). If an expiration date of 1999.365 or 1999.366 is calculated, it will automatically be bumped to 2000.001.

NEWTAPEVOL=
NEWTAPEFILE=

Specifies that the selected entries are to have the volume serial(s) (and optionally the tape file number, up to 4095) of the archive backup file modified. NEWTAPEVOL= will be ignored if TAPEVOL= was not also specified. If the number of volume serials specified (up to 5) is not the same as the number of serials specified by TAPEVOL=, or if NEWTAPEFILE= is specified, then TAPEFILE= must be specified to clearly define which records are to be updated.

REAL
SIMULATE

SIMULATE causes the statement to operate in simulation mode, so that you can verify that the parameters you have specified will have the correct results. It will report on what records it would modify if it was not in simulation mode, but will not update the control file.

REAL makes the desired modifications in the selected records.

The default is SIMULATE. **Innovation strongly recommends that a MODIFY jobstream be run in simulation mode first, to insure that it is modifying the proper records.**

RESET=
SET=

Specifies one or more special indicators to be turned off (RESET) or turned on (SET) in the selected Archive Control File entries. The indicators and their meaning if set are:

AUTOREC – the data set is eligible for Auto-Recall.

DELETE – the data set is flagged for deletion (can also be done by the DELETE and RESET statements).

RESTORED – the data has been restored by ABR.

The default is that the indicators will not be changed.

NOTE: RESET=AUTOREC will not uncatalog the Auto-Recall entry in the system catalog. If this is required, use the RECATALOG statement with SERVICE=UNCATALOG instead.

NOTE: Use SET=AUTOREC if you forgot to specify RECALL=YES when archiving and you now want the data set to be eligible for auto recall. Upon completion of the MODIFY, run the RECATALOG statement to build the Auto-Recall catalog entries.

51.60 CONTINUED . . .

TAPECOPY=	<p>Specifies which archive backup copy is to be modified in records selected. Values are:</p> <p>1 – Only COPY 1 information will be updated.</p> <p>2 – Only COPY 2 information will be updated.</p> <p>BOTH – Both copies (if both exist) will be updated.</p> <p>For expiration date changes (NEWEXDATE and NEWEXPIRE), the default is BOTH. Even if the record was selected because EXPDATE or EXPDAYS matched only one (1) copy, BOTH copies will be changed unless TAPECOPY is specified.</p> <p>For Archive Backup changes (NEWDEV, NEWTAPEVOL, and NEWTAPEFILE), the default is that the change will be made in which ever copy matches on TAPEVOL and, if specified, TAPEFILE. If TAPECOPY is specified for Archive Backup changes, then the specified copy will be changed if and only if TAPEVOL/TAPEFILE matches on that copy.</p>
TAPEFILE=	<p>Specifies a file sequence number (1 to 4095) of the archive tape data set that contains the entry that is to be selected. Ignored unless TAPEVOL= is also specified.</p>
TAPEVOL=	<p>Specifies a volume serial number of the archive tape data set that contains the entry that is to be selected. The volume serial number will be matched against both the COPY 1 and COPY 2 volume serial number lists. Multiple serials (up to 5) can be specified by enclosing the list in parentheses and separate them by commas. However, the complete list of serials will be compared to every record in the Archive Control File. If a given archive run used multiple output tapes, you may have some data sets in a backup file entirely contained on one tape, some entirely on the second tape, and some that are in a file that crossed between the two tapes. You may need multiple MODIFY statements with, for example:</p> <p>TAPEVOL=000001 TAPEVOL=(000001,000002) TAPEVOL=000002</p> <p>to select all the entries on those tapes.</p>
VOL=	<p>Specifies the volume serial number of the disk from which the data sets to be selected were archived. Only those records having a matching originating volume serial number will be selected. Multiple serials may be specified by enclosing them in parentheses separated by commas. Volume serial prefixes can be specified by following the prefix with an asterisk, e.g., VOL=TSO*. See the note on operators above to specify comparisons and ranges.</p>

51.60 CONTINUED . . .

XDSN=

Specifies one or more data set name masks, as defined in [Section 80.14](#). The mask allows for very flexible specification for the data set names to be selected. Multiple masks may be specified in parentheses, separated by commas. Only Archive Control File entries whose name matches one of the masks will be selected. See the note on operators above; XDSN accepts only equal (=) or not equal (≠) comparisons.

Examples: `XDSN=TSO.**.LIST` - all data sets whose first index is TSO and whose last is LIST

`XDSN=ABC+++.**` - all data sets whose first index is ABR plus 3 numeric digits

51.61 FDRARCH RECATALOG STATEMENT

```

RECATALOG      ADATE=yyddd|yyyyddd
                ADAYS=nnnnn

                ,CATERROR=IGNORE|PURGE|RETAIN

                ,CATVOL=CONSTANT|DEFAULT|ORIGINAL

                ,DSN=(dsname,...,dsname)
                ,DSG=(dsgrout,...,dsgrout)
                ,XDSN=(mask,...,mask)

                ,EXPDATE=yyddd|yyyyddd
                ,EXPDAYS=nnnnn

                ,LOG|NOLOG

                ,REAL|SIMULATE

                ,SERVICE=CONVERT|CORRECT|CREATE|MERGE|PERFORM|UNCATLG

                ,VOL=(vvvvvv,.,vvvvvv)

```

**RECATALOG
STATEMENT**

Successful operation of ABR Auto-Recall via the ABR Locate Exit requires that archived data sets remain cataloged with a special indicator in the catalog to indicate that they are archived for Auto-Recall; this field also serves as a pointer to the data set's record in the Archive Control File to avoid lengthy searches; ABR starts at the specified location and reads backwards to find the entry.

If the data set is uncataloged or the DSCBTTR is modified or zeroed, Auto-Recall via the Locate Exit will become impossible. A common cause of this is the IDCAMS function REPRO MERGECAT, which is used to move an ICF catalog, or merge or split catalogs; MERGECAT will zero the DSCBTTR field.

If the location of an entry in the Archive Control File changes due to reorganization (REORG), the DSCBTTR pointer will become increasingly inaccurate and the search for the entry during Auto-Recall may become lengthy.

The RECATALOG statement will attempt to create accurate catalog entries for data sets whose Archive Control File entry indicates that they were archived with the RECALL option for Auto-Recall. RECATALOG will verify that the selected data set is currently not allocated on disk and will then create a new catalog entry if it does not exist, or update the existing entry if the DSCBTTR is zero or inaccurate.

RECATALOG may also be used to uncatalog Auto-Recall entries, and to convert Auto-Recall catalog entries to the volser of MIGRAT (the ABR MIGRAT option) or back to their original volser.

WARNING: Because of the large number of catalog and VTOC accesses required, this may be an expensive and lengthy function. For the SERVICE=PERFORM option (to update the DSCBTTR to improve recall performance), there is also a SORECATALOG statement which uses a SORT to reduce the elapsed time of that function. For details, execute FDRARCH with the parameter "HELP SUB(SORECATALOG)" or browse the FDRARCH member in the FDR ICL (Installation Control Library).

The RECATALOG statement WILL NOT create generation data group catalog entries.

51.61 CONTINUED . . .

RECATALOG STATEMENT The RECATALOG parameters are saved for later processing. Please review [Section 51.50](#) to understand when the Archive Control File will be scanned and the recataloging done.

Many of the operands of RECATALOG will accept comparison operators in addition to a simple equal (=). The operators are:

= or .EQ.	Equal
≠ or .NE.	Not Equal
< or .LT.	Less Than
> or .GT.	Greater Than
<= or .LE.	Less Than or Equal
>= or .GE.	Greater Than or Equal

In most cases, you may also specify an operand more than once, such as:

ADAYS>5,ADAYS<20

to select a range of values.

A record must meet **all** of the selection criteria specified to be selected for recataloging.

OPERANDS	<p>ADATE= Specifies the date of the archive run of the records to be selected. Only data sets archived on this date will be marked. See the note on operators above to specify comparisons and ranges.</p> <p>The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).</p>
	<p>ADAYS= Specifies that data sets that were archived "nnnnn" days ago (from 0 to 32000) are to be selected. See the note on operators above to specify comparisons and ranges.</p>
	<p>CATERROR= Specifies the processing required if an error is encountered while accessing an ICF catalog.</p> <p>IGNORE – The error is to be ignored and treated like a "not-found" condition (i.e., data set is not cataloged) unless the result would be creation of a catalog entry in which case it is treated like CATERROR=RETAIN.</p> <p>PURGE – The error is to cause the entry to be purged from the Archive Control File, if that is an available option. If not, it is treated like CATERROR=IGNORE.</p> <p>RETAIN – Do not continue to process the entry. Retain it for processing after the error within the catalog has been corrected.</p> <p>The default is RETAIN.</p>

51.61 CONTINUED . . .

- CATVOL=** Specifies the volume serial to be used when updating or creating the Auto-Recall catalog entry:
- CONSTANT** – specifies that a constant value ("MIGRAT") is to be used.
- ORIGINAL** – specifies that the original volser that the data set was archived from will be used.
- DEFAULT** – specifies that volser MIGRAT will be used if the MIGRAT option in the FDR Global Option Table is set to YES, otherwise the original volser will be used. The default is DEFAULT.
- Note: if you specify CATVOL= to convert entries to or from the use of MIGRAT, you should also specify SERVICE=CONVERT.
- DSN=** Specifies one or more data set names (up to 44 characters) to be selected. Multiple names may be specified in parentheses, separated by commas. See the note on operators above to specify comparisons and ranges.
- DSG=** Specifies one or more data set name prefixes (up to 44 characters in length) to be selected. Multiple prefixes may be specified in parentheses, separated by commas. Only Archive Control File entries whose name begins with one of the prefixes will be marked. See the note on operators above to specify comparisons and ranges. The XDSN= operand, described below, offers a much more flexible way of selecting data sets.
- Note:** They may be repeated and/or intermingled.
- EXPDATE=** Specifies the expiration date of the archive records to be selected. If both a COPY 1 and COPY 2 backup exist for the data set, the higher of the two expirations is tested. See the note on operators above to specify comparisons and ranges.
- The date is specified as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are considered to be in the 21st century, e.g., 12345=2012.345. For readability, a period may be inserted between the year and day (e.g., 1997.123).
- EXPDAYS=** Specifies that records whose expiration date is within "nnnnn" days of today's date are to be selected. The value can be from 1 to 32000 and can be positive (e.g., EXPDAYS=5 or EXPDAYS=+5 means 5 days from today) or negative (e.g., EXPDAYS=-5 means 5 days ago). It will never consider data sets whose expiration date is 99.000, 99.365 or 99.366 to be expired. See the note on operators above to specify comparisons and ranges.
- LOG** **NOLOG** **LOG** specifies that messages listing the data sets selected are to be printed. **NOLOG** suppresses detail messages.
- The default is LOG unless overridden on a UPDATE statement ([see Section 51.62](#)).

51.61 CONTINUED . . .

REAL **SIMULATE** causes the statement to operate in simulation mode, so that you can verify that the parameters you have specified will have the correct results. It will report on what data sets will be cataloged if it was not in simulation mode, but will not actually catalog them.

REAL catalogs the selected data sets.

The default is **REAL**.

SERVICE= Specifies the services that the RECATALOG statement is to perform for the catalog entries of the selected data sets:

CONVERT – recatalog them for Auto-Recall, changing the volume serial in the catalog as indicated by CATVOL=, if necessary. **CONVERT** can be used to change Auto-Recall entries to and from use of the ABR MIGRAT option.

CORRECT – recatalog them for Auto-Recall if the catalog Auto-Recall indicator is zero.

CREATE – create an Auto-Recall catalog entry for them if not already cataloged. Will not attempt to catalog generations of a GDG.

MERGE – recatalog them for Auto-Recall with the correct Archive Control File pointer if it is incorrect because of a **MERGE** of multiple Archive Control Files.

PERFORM – recatalog them if they are currently cataloged for Auto-Recall but the Archive Control File pointer is more than "VARIATION=" blocks from the actual location. **VARIATION=** is an operand of the **UPDATE** statement.

UNCATLG – uncatalog them (will also reset the Auto-Recall flag in the Archive Control File entry).

All these services verify that the data set is not really on disk before modifying the catalog entry. More than one service can be selected by enclosing them in parenthesis (except **UNCATLG** which must be used alone).

The default is (**CORRECT,CREATE,MERGE,PERFORM**).

VOL= Specifies the volume serial number of the disk from which the data sets to be selected were archived. Only those records having a matching originating volume serial number will be selected. Multiple serials may be specified by enclosing them in parentheses separated by commas. Volume serial prefixes can be specified by following the prefix with an asterisk, e.g., **VOL=TSO***. See the note on operators above to specify comparisons and ranges.

XDSN= Specifies one or more data set name masks, as defined in [Section 80.14](#). The mask allows for very flexible specification for the data set names to be selected. Multiple masks may be specified in parentheses, separated by commas. . Only Archive Control File entries whose name matches one of the masks will be selected. See the note on operators above; **XDSN** accepts only equal (=) or not equal (≠) comparisons.

Examples: **XDSN=TSO.**.LIST** - all data sets whose first index is TSO and whose last is LIST

XDSN=ABC+++.** - all data sets whose first index is ABR plus 3 numeric digits

51.62 FDRARCH UPDATE STATEMENT

UPDATE CNTRESV=nn ,LOG|NOLOG

 ,DISABLE=(option,option,...) ,SEARCH=C|VAF|OBTAIN|NO

 ,ENABLE=(option,option,...) ,VARIATION=nnn

UPDATE STATEMENT The UPDATE statement will execute all of the DELETE, RESET, EXPIRE, MODIFY, and RECATALOG statements that immediately preceded it, making one pass through the Archive Control File and comparing each record to the parameters on those statements, making any required changes to the records that match them.

An implied UPDATE with all defaults taken will be executed whenever one of those statements was entered, and either a FDRARCH statement other than those 5, or the end of the statement input (EOF on SYSIN) is encountered. So the UPDATE statement is not required unless you wish to specify one of the operands.

OPERANDS **CNTRESV=** Specifies the number of Archive Control File blocks (from 1 to 99) to be processed between times when the RESERVE on the volume containing the control file is released. This allows other processors to access the volume at intervals during the reorganization.

The default is 10.

DISABLE= The options specified will be disabled or enabled for this UPDATE operation. Multiple options are enclosed in parentheses and separated by commas. The options and their meaning if enabled are:

ENABLE=

DSNCK validates the data set name of the Archive Control File. It requires that some index level in the name, not the first, is ARCHIVE.

DYNARC dynamically allocates the common Archive Control File named in the ARCDSDN option of the FDR Global Option Table if the ARCHIVE DD statement is not present.

RESERVE protects the Control File with a RESERVE macro against the volume on which it resides.

By default, DSNCK and RESERVE are enabled, DYNARC is disabled, unless overridden in the FDR Global Option Table

LOG Provides a default for the LOG/NOLOG operand on the individual

NOLOG MODIFY and RECATALOG statements that are executed by this UPDATE statement. If LOG or NOLOG is specified on those individual statements, it will take precedence over the operand here on the UPDATE statement. If LOG/NOLOG is not specified on the UPDATE statement nor on the individual statement the default documented for that statement will be used.

51.62 CONTINUED . . .

SEARCH=

Specifies if and how REORG is to verify that data sets flagged for auto recall are accurately cataloged for recall (not on disk).

CVAF uses Common VTOC Access Facility macros.

OBTAIN uses the OBTAIN DADSM macro

NO bypasses this verification.

If CVAF or OBTAIN is specified or defaulted, for every Auto-Recall entry which is being removed from the control file, if the original DASD volume recorded in the control file is online and of the right type (e.g., 3390), REORG will access the VTOC to see if the data set is on that volume. If it is, it will be recataloged to that volume instead of removing the catalog entry.

The default is CVAF.

VARIATION=

If the UPDATE operation includes a RECATALOG statement with SERVICE=PERFORM specified or defaulted, this specifies how many blocks difference (from 1 to 500) there must be between Archive Control File pointer in the Auto-Recall entry in the catalog and the actual location of the entry in the Archive Control File before the catalog is updated with the new location. This is used to avoid expensive catalog updates resulting in minimal recall performance gains.

The default is 5 blocks.

51.63 FDRARCH EXAMPLES

FORMAT AN ARCHIVE CONTROL FILE Allocate an Archive Control File and format it using the default of half-track blocking. On a 3380, at 292 records per track, 500,000 records will require 115 cylinders, so the file will take 2 secondary allocations for a total of 120 cylinders. On a 3390 (348/track), it needs only 95 cylinders so it will format the entire primary allocation of 100 cylinders.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE,UNIT=DISK,DISP=(,CATLG),
//              SPACE=(CYL,(100,10))
//SYSIN         DD     *
                FORMAT RECS=500000
```

FORMAT AN ARCHIVE CONTROL FILE WITH RLSE Allocate an Archive Control File and format it using quarter-track blocking. 300,000 records requires 69 cylinders on a 3380 or 58 on a 3390; since RLSE is specified, FDRARCH will format that many cylinders and release the remainder.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE,UNIT=DISK,DISP=(,CATLG),
//              SPACE=(CYL,300,RLSE)
//SYSIN         DD     *
                FORMAT RECS=300000,B=4
```

BACKUP AN ARCHIVE CONTROL FILE Create a backup copy of an Archive Control File named FDRABR.ARCHIVE.CONTROL on tape. This backup can only be restored by the RESTORE statement of FDRARCH.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE.CONTROL,DISP=SHR
//TAPE1         DD     DSN=COPY.FDRABR.ARCHIVE,
//              UNIT=TAPE,DISP=(,CATLG)
//SYSIN         DD     *
                DUMP
```

RESTORE AN ARCHIVE CONTROL FILE Restore a backup copy of an Archive Control File from a tape created with the DUMP statement of FDRARCH. ENABLE=ERASE is required to allow the restore back on top of an existing control file.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE.CONTROL,DISP=SHR
//TAPE1         DD     DSN=COPY.FDRABR.ARCHIVE,DISP=SHR
//SYSIN         DD     *
                RESTORE ENABLE=ERASE
```

51.63 CONTINUED . . .

**MOVE/
EXPAND AN
ARCHIVE
CONTROL FILE**

Restore a backup copy of an Archive Control File named from a tape created with the DUMP statement of FDRARCH. The output control file is a new file, which can be:

- on a new volume
- on a new device type
- larger than the original

allowing you to move and/or expand the Archive Control File. RESTORE will format the new file (using half-track blocking by default) and restore the records from the backup.

If this control file is to be used for ABR Auto-Recall, the name of the file must be in the ARCDN parameter in the FDR Global Option Table. You must either update the option table, or name this new control file the same as the old one (renaming or deleting the old one first).

NOTE: INNOVATION recommends using FDRARCH rather than FDRDSF or FDRABR to move the Archive Control File to a different device type. FDRARCH will automatically reblock the file for optimum efficiency.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE,
//              DISP=(,CATLG),UNIT=DISK,
//              SPACE=(CYL,200)
//TAPE1         DD     DSN=FDRABR.ARCBACK,DISP=OLD
//SYSIN         DD     *
RESTORE
```

**MERGE
CONTROL
FILES**

Merge three Archive Control Files into one control file. These files are indicated by the DDnames ARCHM1, ARCHM2 and ARCHM3. The new archive file is created using the DD name ARCHIVE; the entire 100 cylinder allocation will be formatted using default half-track blocking.

```
//MERGE        EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE.NEW,
//              DISP=(,CATLG),UNIT=DISK,
//              SPACE=(CYL,100)
//ARCHM1        DD     DSN=FDRABR.ARCHIVE.OLD1,DISP=OLD
//ARCHM2        DD     DSN=FDRABR.ARCHIVE.OLD2,DISP=OLD
//ARCHM3        DD     DSN=FDRABR.ARCHIVE.OLD3,DISP=OLD
//SYSIN         DD     *
MERGE          MERGEDDNAME=(ARCHM1,ARCHM2,ARCHM3)
```

51.63 CONTINUED . . .

**SORT-BASED
REORGANIZA-
TION**

Reorganize the common Archive Control File, using a SORT-based REORG for performance. The sequential backup of the control file is written to a disk GDG; in case of problems it can be restored to its previous condition with the RESTORE statement of FDRARCH.

- ENABLE=DYNARC causes the Archive Control File named in the ARCDN option of the FDR Global Option Table to be dynamically allocated.
- ENABLE=IFNOTCAT causes any data set that was archived for Auto-Recall but which is no longer cataloged to be deleted. This includes old entries for data sets that were archived, were recalled, and were archived again; only the current entry is retained.
- MAXOCCURRENCES=3 retains only the most recent three archived copies of a given data set; others are deleted.
- The defaults of EXPIRE=PURGE, DELETE=PURGE and RESTORE=IGNORE cause expired data sets and those marked for deletion by the DELETE statement of FDRARCH will be deleted during the reorganization. If you want to also delete data sets which have been restored by FDRABR, add this operands:
EXPIRE=PURGE,DELETE=PURGE,RESTORE=PURGE
- SORTALLOC=YES causes all files required for sorting to be dynamically allocated (if your SORT product has different needs, you may need to provide necessary DD statements).

Any deleted entries still cataloged for Auto-Recall will be uncataloged and Archive Backup files which are no longer referenced after the REORG will be deleted and uncataloged.

Innovation strongly recommends that the SORT-based REORG be used for most reorganizations.

```
//REORG      EXEC  PGM=FDRARCH,REGION=0M
//SYSPRINT   DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//TAPE1      DD    DSN=OPNS.ARCHIVE.BACKUP(+1),
//              UNIT=DISK,SPACE=(CYL,(100,25),RLSE),DISP=(,CATLG)
//SYSIN      DD    *
              REORG SORT=YES,SORTALLOC=YES,ENABLE=(DYNARC,IFNOTCAT),
              MAXOCCURRENCES=3,DUMPDEVICE=DISK
```

**DELETE AND
REORGANIZ-
ATION**

Mark specified data sets for deletion from the Archive Control File, then reorganize that file to actually perform the deletion. The DELETE statement will first scan the entire control file, marking the data sets whose name matches the mask specified and which were archived more than 30 days ago. Then the REORG takes a sequential backup of the control file to a tape and reads it back to reload the Archive Control File, excluding the flagged records. Any deleted entries still cataloged for Auto-Recall will be uncataloged and Archive Backup data sets which are no longer referenced after the REORG will be deleted and uncataloged.

Note: the SORT-based reorganization shown in the previous example is recommended for most REORGs since it is much more efficient when catalog accesses are required. The non-SORT REORG shown here can be used for small Archive Control Files (typically under 100,000 records) or when a limited number of records are affected as shown in this example.

```
//REORG      EXEC  PGM=FDRARCH,REGION=0M
//SYSPRINT   DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//ARCHIVE    DD    DSN=FDRABR.ARCHIVE,DISP=SHR
//TAPE1      DD    DSN=TECH.REORG.ARCHIVE,
//              UNIT=CART,DISP=(,CATLG)
//SYSIN      DD    *
              DELETE ADAYS>30,XDSN=TECH.*.LIST*
              REORG DELETE=PURGE
```

51.63 CONTINUED . . .

**SIMULATE
REORGANIZ-
ATION**

Simulate the reorganization of the production Archive Control File. This allows you to verify the correct operation of your parameters and to see the potential results. This will use the SORT-based REORG. Since the TEMPARCH DD statement is provided, the reorganized control file is written to that data set (which must be as large as the input control file); you can then run reports against it to verify correct operation.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR,DSN=FDRABR.ARCHIVE
//TEMPARCH      DD     DISP=(,CATLG,DELETE),DSN=GEORGE.ARCHIVE,
//              UNIT=SYSDA,SPACE=(CYL,50,RLSE)
//SYSIN         DD     *
REORG SIMULATE,SORT=YES,SORTALLOC=YES,ENABLE=IFNOTCAT,
      MAXGENERATION=6,MAXOCCURRENCES=1,
      DELETE=PURGE,EXPIRE=PURGE
```

**MARK FOR
DELETION**

Mark for deletion all archive records created in January 1997 from TSO volumes, but reset the delete indicator for certain data sets. Result will be that all records except those with a first index of SAVEDSN will be deleted during the next reorganization. The UPDATE statement actually executes the DELETE and RESET statements in one pass of the control file. ENABLE=DYNARC causes the Archive Control File named in the FDR Global Option Table to be dynamically allocated for this operation.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
DELETE         ADATE.GE.97001,ADATE.LE.97031,VOL=TSO*
RESET          ADATE.GE.97001,ADATE.LE.97031,VOL=TSO*,DSG=SAVEDSN.
UPDATE         ENABLE=DYNARC
```

**EXTEND
EXPIRATIONS**

Extend by 30 days the expiration dates of all data sets which were archived in 1997. Also set the expiration date of all data sets archived to disk volume ARC003 to 14 days from today. For backups on tape, you may need to manually extend the expiration of the tapes involved through a tape management system utility. The UPDATE statement actually executes the MODIFY statements in one pass of the control file. ENABLE=DYNARC causes the Archive Control File named in the FDR Global Option Table to be dynamically allocated for this operation. LOG cause all modified data sets to be listed. Omit the REAL operands to perform the modification in simulation mode; REAL is required to actually modify the records.

```
//FDRARCH      EXEC   PGM=FDRARCH
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//SYSIN         DD     *
MODIFY         ADATE.GE.97001,ADATE.LE.97366,REAL,
              NEWEXPIRE=EXPIRE,DAYS+=30,
MODIFY         TAPEVOL=ARC003,NEWEXPIRE=TODAY,DAYS+=14,REAL
UPDATE         ENABLE=DYNARC,LOG
```

**CORRECT
AUTO-RECALL
CATALOGS**

An IDCAMS "REPRO MERGECAT" operation has been used to move the system catalog entries for data sets beginning with ABC and XYZ to a new catalog; this operation zeros the DSCBTTR field in the catalog entry making Auto-Recall of those data sets impossible. The RECATALOG statement will find all entries in the Archive Control File for those names which are flagged for Auto-Recall, and correct the DSCBTTR fields in the catalog to re-enable Auto-Recall for them.

```
//STEP1        EXEC   PGM=FDRARCH,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=FDRABR.ARCHIVE,DISP=SHR
//SYSIN         DD     *
RECATALOG      XDSN=(ABC.***,XYZ.**),SERVICE=CORRECT
```

51.63 CONTINUED . . .

**MAINTAIN
AUTO-RECALL
CATALOGS**

This RECATALOG will scan the entire Archive Control File and insure that the Auto-Recall entries in the system catalogs for all data sets flagged for Auto-Recall are correct. It will create catalog entries for any that are not currently cataloged and correct any that have an improper or zero Auto-Recall pointer. The Auto-Recall pointer in the catalog points to the record in the control file where ABR is to start searching for a data set to be recalled; if the location of a given data set has moved more than 10 blocks due to REORG, the catalog pointer will be updated to improve recall performance. Do not display the data sets that were recataloged (NOLOG).

```
//STEP1      EXEC   PGM=FDRARCH,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
  RECATALOG   NOLOG
  UPDATE     ENABLE=DYNARC,VARIATION=10
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

52.01 FDRAPPL APPLICATION BACKUP

ABR Program FDRABR (Automatic Backup Restore) automates the execution of FDR full-volume and data set backups and restores for the purposes of:

- **data availability** - creating backup copies of DASD data sets to protect against physical loss or logical damage.
- **space management** - identifying data sets which do not need to be on DASD (usually based on the last date they were used) and moving them to backups (called "archiving"), from which they can be automatically recalled if needed. Data sets that will never be needed again can be scratched without creating a backup.
- **disaster recovery** - creating backups from which all or part of your DASD data can be quickly recreated at a disaster recovery site.
- **tape mount management (TMM)** - maximizing tape volume usage by staging small tape data sets to a disk buffer, then moving many of them to a tape volume.

These objectives are accomplished by the use of several ABR functions:

- This section describes Application Backup, also known as FDRAPPL, used for "data availability" and "disaster recovery" at the application level. FDRAPPL is automatically included if you are licensed for ABR, but it can also be licensed as an option to FDR.
- ABR Volume Backups (DUMP TYPE=FDR/ABR/AUTO/DSF) are found in [Section 50](#).
- ABR ARCHIVE (DUMP TYPE=ARC) and SUPERSCRATCH (DUMP TYPE=SCR) are in [Section 51](#).

ABR BACKUPS All the backup functions of ABR share these common characteristics:

- the backups are in standard full-volume (FDR) or data set (DSF) backup format, described in [Section 02.02](#). If necessary, data can be restored from ABR backups with FDR or FDRDSF, but ABR automates the restore process.
- although ABR backups require a separate backup data set for each DASD volume processed, ABR will automatically stack multiple backup data sets on tape, creating multi-file tapes, to make best use of today's high-capacity tape volumes (such as IBM Magstar and StorageTek Redwood and 9840). If necessary, multiple output tape volumes are used. No special JCL is required since ABR will handle the file creation internally.
- ABR backups to disk are also supported. ABR will automatically allocate the backup data sets on the backup volumes you specify. Disk backups are usually not used with Application Backup.
- the output devices (tape and/or disk) are specified in the ABR batch JCL. However, you only need to identify the output device; ABR will automatically name and create every backup data set and catalog it if required.
- You can specify in the JCL that two copies of each backup are to be created, even though the input disk is read only once. In ABR these are known as COPY 1 and COPY 2 ([see Section 52.04](#) for details). The ABR utilities FDRTCOPY and FDRTSEL ([Section 60](#)) can be used to create COPY 2 from COPY 1, but there are some limitations for Application Backup.
- If multiple output devices are specified in the JCL, ABR will automatically use internal subtasking to process more than one DASD volume concurrently.

52.01 CONTINUED**FDRAPPL
APPLICATION
BACKUP**

FDRAPPL Application Backup, as its name implies, is designed to backup all the data sets that relate to a given application, from whatever disk volumes they reside. Since application data sets are usually cataloged, you can have FDRAPPL search the MVS catalogs for data sets matching one or more name masks, quickly and easily selecting the data sets and their cataloged volumes to be processed.

FDRAPPL will process all of the data sets selected on the volumes selected. Since FDRAPPL always creates one backup data set per disk volume, it will automatically stack multiple backup files on the output tape, depending on how many disk volumes were read. Usually a single tape is used for output, creating one tape with all the selected data (a duplicate backup can be created).

FDRAPPL does use a data base to record the data sets that were backed up. This data base is called an Application Control File (ACF); it has the same format as the Archive Control File described in [Section 51.01](#). The ACF is very compact, recording several hundred data sets in a single track. A separate ACF will be used for each application.

If data sets must be restored from an Application Backup, the information on their location is read from the ACF to build a list of backup data sets to be dynamically allocated and read. You can select the data sets to be restored, or simply restore the most recent copy of every data set in the ACF. Obviously, access to the ACF is a critical requirement for automated restore from FDRAPPL. It is possible to do manual FDRDSF restores, but it requires a TAPEX DD for every backup file on the backup tape.

There are several ways of managing the ACF. Details are found in [Section 52.10](#).

Many application jobstreams contain a series of specialized backup steps in order to create backups of application-oriented data sets to provide restart and recovery capabilities just for that application. These backups (often IEBGENER or IDCAMS REPRO steps) may be very time-consuming and may use many tapes. FDRAPPL may be used as a single-step high-speed replacement for those specialized backups. If copies of Application Backup tapes are sent to offsite storage, they may also be used to recover application data sets at a disaster recovery site.

52.01 CONTINUED**WHY
FDRAPPL?**

Why would you want to use FDRAPPL instead of relying on ABR Volume Backups?

Volume Backups are usually taken at a fixed time each night. For many applications, this is the wrong time and does not provide the recovery and restartability they need. Application Backups can be inserted into the application jobstreams, to provide backup at a point in processing chosen by the application developers.

Frequency and retention of the Application Backups is under control of the application, instead of the Data Center.

At a disaster recovery site, critical applications can be recovered first and less urgent applications at a later time.

Under SMS and other disk pooling software, users often are unaware of the disk volume serials of their data sets. FDRAPPL allows you to provide one TAPE1 DD (with optional duplicate backup TAPE11), and to select data sets from the system catalogs without being concerned about the volumes they reside on.

Application recovery uses data set (FDRDSF-type) restores, so the data sets may be directed to a different set of volumes than those they were dumped from, and may be spread out on fewer or more volumes than they originally occupied. Remember that multi-volume data sets must be restored to the same number of volumes they were dumped from, with the original amount of space on each.

For recovery at your own data center, FDRAPPL is useful if a critical application must have recovery, owns data sets which reside on multiple disk volumes, and the backups cannot be scheduled as part of the normal Volume Backups. The Application Backup is usually done after the batch or online processing for the application has completed and updates to the data sets are quiesced. Since only data sets from this application are dumped, using all of the performance of an FDR backup, the down time for the application is reduced to a minimum. The application backup can be executed as a step in the actual application jobstream, so that it will be executed at the point that is required by that application.

FDRAPPL is comparable to ABARS under IBM's DFHSM or DFSMSHsm products.. Compared to ABARS, FDRAPPL:

- a. is significantly faster,
- b. offers more flexible selection criteria.

**IMPLEMEN-
TATION**

There are several different ways that you can implement FDRAPPL, each with its own advantages. You must choose the one that best meets your needs; the choice may depend on whether you intend to use FDRAPPL to recover applications at a disaster site, or at your prime site, or both.

The differences in the techniques are primarily in the usage of the Application Control File. Remember that FDRAPPL records the Application Backups in a Application Control File, and that Application Control File is required to restore the data sets. The two major techniques are:

- One permanent Application Control File for each application.
- The Application Control File for each application is a GDG, creating a new generation for each backup step.

See [Section 52.10](#) for more details on these techniques, including specific examples.

SIMULATION

All functions of FDRAPPL can be run in simulation mode, allowing you to verify that the correct data will be selected when run for real.

CONTINUED . . .

52.01 CONTINUED

ABR UTILITIES ABR includes a utility program, FDRARCH, for maintaining the Application Control File. If you choose to maintain a permanent ACF for each application, you will need periodically run the REORG function of FDRARCH to purge those obsolete entries and make room for new ones. FDRARCH is documented in [Section 51.50](#).

FDRABRP ([Section 53](#)) can be used for simple reporting on data sets recorded in an ACF, and FDREPORT ([Section 54](#)) can be used for more sophisticated reporting. The ABR ISPF dialog called SRS (Search, Report, and Services) can also display information on Application Backup data; it is also in [Section 54](#). In each case you must point to the proper Application Control File to be used as input.

**TAPE FORMAT
AND NAMING
CONVEN-
TIONS**

As indicated earlier, the backup files created by FDRAPPL are in standard FDRDSF (data set backup) format. In this format, each backup file can contain data only from one DASD volume. So, if multiple DASD volumes are to be processed, FDRAPPL must create multiple backup files on the output tape or disk. Unless you select data sets from only a single disk volume in an FDRAPPL step, FDRAPPL will **always** create multiple backup files.

Since FDRAPPL must be able to uniquely name each tape file, and must be able to record the backup file in a way that it can easily be retrieved, FDRAPPL uses a special naming convention for the Application Backup files. The name contains the disk volume serial, the date of the backup and a uniqueness character in case data is archived multiple times per day, so each FDRAPPL Backup from a given disk volume will have a unique name. This is the same convention used by Archive Backups ([Section 51](#)), except for the hi-level index.

The format is: **userindx.Vvvvvvv.bnyydddx**

where:

userindx	Is the same as the hi-level index of the Application Control File, so it is a hi-level index chosen by the user. It is usually an index belonging to the application owning the backup.
vvvvvv	is the volume serial of the disk volume from which the data sets were backed up. FDRAPPL creates one backup file for each disk volume processed in a given ABR run.
n	is the copy number. FDRAPPL always creates COPY 1 and can optionally create COPY 2.
yyddd b and x	Is the Julian date of the backup job (5 digits) are qualifiers added to make the name unique if multiple Application Backups are created for the same disk volume on the same Julian date. However, they are used only if the name FDRAPPL is trying to create is already cataloged. Since FDRAPPL data sets are usually not cataloged, "b" and "x" usually have their default values of B and A (see Section 51.01 for details)

Example: PAYROLL.VTSO002.B197147A (COPY 1 created on 97.147 for volume TSO002)

Warning: you cannot override this naming convention. As you will see in [Section 52.04](#) you must specify a data set name on the output (TAPEX) DD statements to satisfy MVS requirements, but that name will be ignored and the ABR-generated name used in its place.

52.01 CONTINUED

The backup files created by FDRAPPL are usually not cataloged. For most such backups, all the information required is stored in the data set records in the Application Control File. When restoring, information about the backup file to be read is obtained from the ACF so no entry in the catalog is required. If you like, you can override this: ARCCAT=ALL forces FDRAPPL to catalog all backup files, while ARCCAT=NORMAL follows the rules for Archive Backups in [Section 51.01](#).

By default, FDRAPPL will backup the Application Control File itself as the last file on the backup tape created on both TAPEX and TAPEXX using a FDRDSF-type backup of just that one data set. If your step included more than one TAPEX DD, the ACF is backed up to the last tape that was active. This makes the tape self-contained, with all of the requested data plus a copy of the ACF required to restore from the tape. By default, the data set name of the backup file itself will be the same as the name of the ACF, except that the index level of ARCHIVE will be replaced by ARCBKUP (on TAPEX) and ARCBKU2 (on TAPEXX). However, you can control the backup data set names with the ARCB1DSN= and ARCB2DSN= operands; this is recommended.

**BACKUP
RETENTION
AND TAPE
MANAGE-
MENT**

Depending on the requirements for retention of backups established by each application, you may choose to control the retention of Application Backups in several ways:

- You can let your tape management system expire backups based on a retention period (RETPD=). This is known as "date control". Both your tape management system and the Application Control File will have the same expiration date for the data created by a given backup. The tape management system will scratch the tapes on that day. This is usually used with a permanent Application Control File ([see Section 52.11](#)).
- You can retain backups based on the number of times the backups have been run. To do this the Application Control File (or at least the backup of the ACF) must be setup as a GDG ([see Section 52.12](#)). You must specify that your tape management system is to use "catalog control" for these tapes (specified by EXPDT=99000 for many tape management systems). Since the ACF backup is the only cataloged file on the tape (or multi-volume tape set), the entire tape set will be retained until that file is uncataloged. It will eventually be uncataloged as new generations are created (new backup runs) and the tape management system will scratch the tapes.
- If you have no tape management system, you must establish the necessary manual procedures for identifying and expiring FDRAPPL backups.

Note: if you have a tape management system from any software vendor, you should enable the TMS (Tape Management System) option in the FDR Global Option Table ([see Section 90](#)). The TMS option changes slightly the way that FDRAPPL handles files on tape to be compatible with restrictions of some such systems.

CONTINUED . . .

52.01 CONTINUED**BACKUP
RETENTION
AND TAPE
MANAGE-
MENT**

When an Application Backup is taken, an expiration date is assigned to each backup file. It can come from several sources, in this order of priority:

- if the RETPD= operand is specified on the FDRAPPL DUMP statement, FDRAPPL uses that to calculate an expiration date which will be assigned to all COPY1 (TAPEX) backups created in that FDRAPPL step. If the RETPD= operand is not specified but COPY2 (TAPEXX) backups are also created, the same expiration is assigned to those backups.
- If the RETPD2= operand is specified on the DUMP statement, FDRAPPL uses that retention period to calculate an expiration date which will be assigned to all COPY2 (TAPEXX) backups created in that FDRAPPL step.
- if the TAPEX or TAPEXX DD statement pointing to the backup tape contains the EXPDT= operand, the date specified is assigned to all backups written to that DD. For some tape management systems this value can also be a keyword; for example, EXPDT=99000 indicates "catalog control" to many tape management systems.
- If the DD statement contains the RETPD= operand, MVS uses that to calculate a expiration date, which FDRAPPL will assign to each backup written to that DD.
- If none of the above RETPD/EXPDT operands are given, the default retention period of 365 days (1 year) is used to calculate an expiration date which will be assigned to each backup.

By whatever means it is calculated, the expiration date is recorded in the Application Control File record for each data set included in that Application Backup file. If both COPY 1 and COPY 2 of a backup are being created, the expiration of each copy is recorded separately in case they are different.

The expiration date will be recorded in the tape labels of the backup file and will also be recorded by your tape management system, if you have one. If the expiration date is a real date (not a keyword such as 99000) then your tape management system will probably return the tape to scratch status on that date.

If you create both COPY 1 and COPY 2 of your backups, you can use one copy for onsite restores and the other copy for offsite restores at a disaster/recovery site. Most tape management systems include vaulting support, allowing you to select tapes to be sent offsite and returned when no longer required. Most users send COPY 2 offsite. Since the data set name used by FDRAPPL includes the copy number, you should be able to easily send COPY 2 backups offsite while retaining COPY 1 onsite.

FDRAPPL

To perform Application Backups, you must create FDRAPPL steps and insert them into appropriate application jobstreams. You may want to create standard FDRAPPL steps or cataloged procedures (procs) which can be invoked by application programmers. [Section 52.11](#) and [52.12](#) contains examples of such steps which you can customize for your installation; all of the examples shown in this manual are also in the JCL library loaded as part of FDR's installation ([Section 90](#)).

Recommended techniques for executing Application Backups are found starting in [Section 52.10](#).

You will probably also want to establish procedures to allow application programmer to create jobs to restore all (o a subset) of the data sets in a given Application Backup. Restore techniques are also detailed starting in [Section 52.10](#).

52.02 DATA SET SELECTION

FDRAPPL The rules for data set selection in FDRAPPL are very simple:

- **SELECT CATDSN=** statements are normally used to select the data sets to be included. The CATDSN= value can be a single data set

```
CATDSN=FINANCE.LOADLIB
```

a very simple mask

```
CATDSN=PAYROLL.**
```

for all data sets starting with PAYROLL

or a complex mask

```
CATDSN=*.SUB%%*.*.DAT
```

The system catalogs are scanned for the indicated data sets. All such data sets found in the catalogs will be added to FDRAPPL's data set list, and the volumes to which they are cataloged will be added to FDRAPPL's list of volumes to be processed.

- **SELECT DSN=** statements may be used to select uncataloged data sets. They can also be used avoid the overhead of the catalog search by going directly to specified disk volumes. The VOL= or VOLG= operands specify the volumes on which to search for these data sets; those volumes are automatically included in FDRAPPL's list of volumes to process. DSN= accepts the same values as CATDSN=.

- If you need to exclude certain data sets from the backup, an **EXCLUDE DSN=** statement can be used. **EXCLUDE** statements must precede **SELECT** statements.

ICF VSAM FDRAPPL can only backup ICF VSAM files by cluster name. You cannot select individual components. Once FDRAPPL selects a cluster from a volume, all components of the cluster (including alternate index components) are dumped.

52.02 CONTINUED**RESTORE
FROM
APPLICATION
BACKUP**

To restore a data set from Application Backup, you must execute an ABR "RESTORE TYPE=APPL" job and provide one or more SELECT statements to identify the data sets required. You can use SELECT ALLDSN to restore all data set recorded in the Application Control File.

The JCL must also point to the Application Control File (ACF) in which the backups were recorded; if this ACF is not already on disk, you will need to execute a FDRDSF step to restore it from the backup at the end of the backup tape.

Since the ACF is organized in chronological order based on backup date, FDRAPPL will search for the data sets requested by reading it in reverse order. If you use a permanent Application Control File and a data set has been backed up more than once, this will find the most recently archived version. If you want an earlier version, you can specify ADATE= (date of backup) or OLDBACKUP= (relative version) to tell FDRAPPL which version to restore.

If a multi-volume data set was backed up, each volume will have a separate entry in the ACF, but each will be flagged to indicate it was part of a multi-volume entry. FDRAPPL will continue to search the ACF until it finds all of the entries belonging to the data set.

The restore process is identical to that of FDRDSF. Please review [Section 20.01](#) for details on allocating, restoring, and cataloging data sets during the restore. [Section 52.08](#) contains details on output volume selection for FDRAPPL restores.

FDRAPPL will default to restoring from COPY 1 of each backup required unless COPY 1 was not created or has expired. COPY2 will then be selected if it exists. If TAPE=EXP is specified, FDRAPPL will ignore the expiration date test and attempt to restore from COPY 1 even if it is expired.

52.03 PROCESSING OPTIONS AND REQUIREMENTS**FDRAPPL
OPERATIONS**

The first control statement in the FDRAPPL input must be a DUMP or RESTORE statement. It may be followed by SELECT and optionally EXCLUDE statements to specify data sets which must be backed up. FDRAPPL will accept up to 100 SELECT/EXCLUDE statements in a single execution, unless that limit is overridden by MAXCARDS=. Note that if a SELECT CATDSN= is translated into an internally-generated SELECT for each data set selected from the system catalogs, so you may need to specify a large MAXCARDS= value if many data set will be selected from the catalog.

**FDRAPPL
OPTIONS**

The DUMP statement contains a TYPE=APPL operand to request FDRAPPL.

Note to users of previous releases: prior to FDR/ABR V5.2, DUMP TYPE=ARC,SCRATCH=NO was used to invoke Application Backup. Existing jobstreams using SCRATCH=NO will continue to work. However TYPE=APPL invokes some special options for Application Backup and is recommended. SCRATCH=NO may not work exactly as described in this section.

If the FDRAPPL step JCL contains a single TAPEX DD statement, FDRAPPL will select the first disk volume on its list and backup all selected data sets on that volume, creating file 1 on that output tape. When complete, it selects the next volume and creates file 2 on the tape, etc., until all volumes are processed.

However, you may have up to 9 TAPEX DDs in the Archive step. If you have more than one, FDRAPPL will select the first "n" volumes on its list and assign them to the "n" tape drives, executing the backups in parallel with internal subtasking, and creating file 1 on each output tape. As each backup completes, FDRAPPL selects the next disk volume and assigns it to the tape which just became idle, creating the next file number on that tape. Remember that the backup of the Application Control File will be placed only on the end of the last TAPEX which was active; this may cause problems if you are depending on that backup being on the tape for tape management catalog control (see [Section 52.01](#)). See "Memory Requirements" below for limitations on the use of multiple TAPEX DDs.

**MEMORY
REQUIRE-
MENTS**

FDRAPPL has the same memory requirements as ABR. See "Memory Requirements" in [Section 50.03](#).

52.03 CONTINUED

COMPRESS OPTION FDRAPPL can be instructed to compress the data on the sequential backup file using Innovation's own proprietary software compression algorithm.

It is **not recommended** for backups to tape attached by ESCON or FICON channels because of the speed of the channel. All ESCON/FICON-attached tape drives include hardware compression which will provide savings similar to that of FDR compression. Software compression will be ignored for backups created with the HSDM disk hardware option (DCT=YES) since the data is pre-compressed.

NOTE: all FDR restores will automatically recognize a compressed backup file and decompress it. No special option is required to restore a compressed backup.

DUPLICATE TAPE OPTION FDRAPPL has an option to create a duplicate or second copy of the backup tape during dump processing. The primary copy is called COPY 1 and the duplicate is called COPY 2. The copy number becomes part of the name of the backup data set as described in [Section 52.01](#).

While dumping a disk to a TAPE DD statement, the duplicate backup will be written to the TAPE DD statement (same "x" value twice) if it is present. You may have TAPE DDs for some TAPE DDs and not for others in the same step but this is not recommended with FDRAPPL since you cannot predict which disks will be written to which tapes.

Memory requirements do not increase with the use of the duplicate tape option.

STORAGETEK EXHPDM SUPPORT FDRAPPL supports the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM takes multiple concurrent tape outputs (such as FDR backup TAPE or TAPE DD statements) and directs them to a smaller number of tape drives, interleaving the data in a single tape file. ExHPDM is invoked by adding the SUBSYS= operand to the TAPE DD statements. [See Section 80.33](#) for more details.

SECURITY Complete details on the security options of the FDR system are found in [Section 80.15](#) "Security".

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12](#) "Security Options".

If your security system is RACF, or another security system which supports the SAF (Security Authorization Facility) interface, such as ACF2 or TOP SECRET, you can enable the ALLCALL option. For FDRAPPL this results in these security checks:

- for Application Backup (DUMP TYPE=APPL), ABR will check to see if your userid has at least READ authority to the entire input volume; under RACF this means that you are authorized to the input volume serial under the DASDVOL security class (other security systems have similar ways of defining volume authority). If you do have this volume authority, no additional checks are done on that input volume. If you do not have volume authority or the input volume is not protected by your security system, then ABR will check if you have at least READ authority under the DATASET security class to every data set selected from the input volume. Any data sets to which you are not authorized will be bypassed with an error message. Since Application Backups will be run under a security userid belonging to an application, the DATA SET security checks will normally be used.
- for data set restores (RESTORE TYPE=APPL), ABR will check if you have at least UPDATE authority under the DATASET security class to every data set restored. Any data sets to which you are not authorized will be bypassed with an error message. If an output data set must be allocated, the operating system will check if you have CREATE/ALLOCATE authority for the data set (this is done even if ALLCALL is not enabled). The user may need READ authority to the Archive Backup files on tape or disk (even if ALLCALL is not enabled).

52.03 CONTINUED

**DATA SET
ENQUEUE
OPTION**

You can request, via the DSNENQ= operand, that each data set being backed up or restored be tested to see if it is in use. A data set is considered in use if any job or TSO user has a DD statement or dynamic allocation for that data set name.

In-use data sets are tested by doing an exclusive ENQ with a major name of SYSDSN and a minor name of the data set name itself, for each selected data set found in the VTOC of the input disk; this resource will be enqueued by any other task allocating the data set so our ENQ will fail if it is in use. Note that FDR cannot tell if the data set is being used for input or output. It also cannot tell what volume an active data set is on, so FDR will think a data set on one volume is active even if a data set by the same name on another volume is really the active one; these are MVS limitations.

Optionally you can request that inactive data sets be enqueued to FDRAPPL during the backup or restore, to insure that no other job or TSO user can access the data set until FDRAPPL is done.

For backups, in-use data sets will still be dumped by default, but you must be aware that the backups of data sets which are being updated during the backup may be unusable, depending on the nature and format of the data. If you wish to bypass the backup of active data sets during an Application Backup, specify the ENQERR=BYPASS operand.

For restores, FDRAPPL will attempt to enqueue any data sets that it allocates on the output disks, to insure that no other task tries to use them until the restore is complete, but if the ENQ fails, the data set is still restored. But for existing data sets, if the ENQ fails, the restore will be bypassed.

The DSNENQ= operand has 4 possible values:

USE — data sets will be enqueued for the duration of the backup or restore on this disk volume. For data sets that are active, a FDR158 warning message is issued and the data set is not enqueued. This is the most frequently used option. **This is the default for TYPE=APPL restores.**

TEST — data sets will only be tested to see if they are enqueued to another task at the time that the backup or restore on this volume starts. For data sets that are active, a FDR158 warning message is issued. The data set will not be enqueued and other tasks may enqueue it and possibly update it while the backup or restore is proceeding.

HAVE — The data sets will be enqueued for the duration of the backup or restore. If a data set is in use, the MVS operator must interact with FDRAPPL to decide how to proceed; a message (FDRW27) is issued to the MVS console, and the operator can respond:

WAIT - wait for the data set to become available; if it is not eventually dequeued, the FDRAPPL job may time out, so the operator must know which data sets are in use by long-running jobs or tasks.

NOWAIT - do not enqueue the data set. The FDR158 warning message is issued.

RETRY - try the enqueue again. If it fails again, the FDRW27 message is reissued.

NONE — No data set ENQ will be issued. **This is the default for backups.**

NOTE: If a data set name appears in a DD statement with DISP=SHR within the FDRAPPL job (not necessarily in the FDRAPPL step), and you specify DSNENQ=USE or TEST, FDRAPPL will change the scheduler enqueue for the data set to EXCLUSIVE (DISP=OLD). The data set may be unavailable to other tasks until the FDRAPPL job ends.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Recommendation: Use DSNENQ= to prevent other tasks from updating or reading data sets being dumped or restored. However, users of FDRAPPL may be able to be sure that data sets may be safely backed up, especially if the application jobstream is the only user of the data sets, so DSNENQ=NONE may be appropriate. Member ENQ in the FDR ICL (Installation Control Library) has more information on data set ENQs.

If HFS=QUIESCE is specified, special processing is done for HFS data sets (Hierarchical File System, used by OS/390 Unix System Services). If the SYSDSN ENQ cannot be acquired, this may mean that the file system is mounted to OS/390 Unix, so FDR will attempt to quiesce the file system during the backup. Details on the quiesce function are found in [Section 80.11](#).

CONTINUED . . .

52.03 CONTINUED**STEP
TERMINATION**

If no errors occur during the execution of FDRAPPL, the FDRAPPL jobstep will end with condition code 0 (zero).

If errors do occur, they are generally indicated by a error message; occasionally they are indicated only by a user ABEND (Uxxxx). Depending on the nature of the error, the step may end one of several ways:

Some errors are critical. The jobstep ends immediately with a user ABEND.

Some errors are critical only to a particular operation. For example, during a backup, some errors cause the backup of a particular disk to terminate immediately, but FDRAPPL may continue and attempt to backup other disks requested in the same step.

Some errors are non-critical and the messages are warnings only. FDRAPPL will complete the current operation.

For the last 2 conditions above, a flag is set indicating that a non-terminating error occurred. At step termination, it tests the flag; if it is on, the step will terminate with return code 12 to call your attention to the errors. Remember that RC=12 indicates that some or all of the functions you requested **did complete** but you must examine the error messages to determine the impact of the errors.

If you prefer to get a different return code or a U0888 abend on a non-terminating error, the ABRCC option in the FDR Global Option Table can change it to a non-zero return code of your choice or ABEND ([see Section 90](#)).

**DYNAMIC
ALLOCATION**

FDRAPPL will dynamically allocate disk volumes as needed. As long as the required volumes are online, FDRAPPL can dump from and restore to any required disk volumes.

For restore operations, FDRAPPL will dynamically allocate each required backup tape if the DYNTAPE option is specified. If required backups are on mixed device types (such as 3480s and 3490Es), FDRAPPL will automatically mount each tape on the proper device type. For automated tape libraries, a drive in the proper library will be allocated. If you have multiple tape libraries, you may need to enable the DYNDEALC option in the FDR Global Option Table (ISPF panel A.I.4.4, see [Section 90](#)).

For data set restores, FDRAPPL will sort the list of backup files required and mount the backup tapes in an order which minimizes the amount of tape movement required.

52.04 FDRAPPL BACKUP JOB CONTROL REQUIREMENTS

The following Job Control Statements are required to perform FDRAPPL backups..

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRABR), region requirement (REGION=, [see Section 52.03](#)), and optional PARM= operand.

If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid DUMP statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRABR,PARM='DUMP TYPE=APPL,RETPD=10'
//FDR EXEC PGM=FDRABR,PARM='DUMP TYPE=APPL/ SELECT CATDSN=A.B.C'
```

If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

SYSPRINT DD STATEMENT Specifies the primary output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

SYSPRINx DD STATEMENT Specifies the secondary output data set for messages related to the matching TAPEX DD statement. A SYSPRINx is required for each TAPEX present in the step. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

ARCHIVE DD STATEMENT Specifies the Application Control File in which all data sets dumped by this FDRAPPL step will be recorded.

It may reference a permanent Application Control File (unique for each application) in which case entries will be accumulated from multiple FDRAPPL runs. Specify DISP=SHR since FDRAPPL internally serializes access. For example,

```
//ARCHIVE DD DSN=INVOICE.BACKUP.CONTROL,DISP=SHR
```

If the ARCHIVE DD references a new data set, FDRAPPL will format it as data sets are recorded in it. Since each track of the control file can record the backups of several hundred data sets, a few tracks are usually all that is required. For example,

```
//ARCHIVE DD DSN=PAYROLL.ARCH1(+1),UNIT=DISK,
// SPACE=(TRK,(3,1),RLSE),DISP=(,CATLG)
```

If ARCBKUP=DSF is specified or defaulted, and ARCB1DSN (and ARCB2DSN, if applicable) are not specified, then the data set name of the Control File must include the index level 'ARCHIVE'; it will be replaced by ARCBKUP and ARCBKU2 to create the name of the backup files containing the Application Control File.

For SIMULATION runs, this DD statement is not required.

52.04 CONTINUED

**TAPE1 DD
STATEMENT**

Specifies the COPY 1 backup tape to be used for FDRAPPL. If data sets from multiple disk volumes are dumped, FDRAPPL will create multiple files on the tape (or tape aggregate if more than one tape volume is used), one file for each disk volume.

Usually only TAPE1 is required for FDRAPPL, and FDRAPPL will process disk volumes one at a time. If a large amount of data from multiple disks is to be dumped, you may specify additional TAPEX DDs and FDRAPPL will process up to 9 disks in parallel (however, region requirements may limit the number of concurrent subtasks; see "Memory Requirements" in [Section 52.03](#)). There must be a SYSPRINX DD and optionally a TAPEXX DD matching each TAPEX DD. Do not specify DSNENQ= if multi-volume data sets might be involved.

On each TAPE DD you should specify:

DSN=	—a data set name is required by MVS, but it will be overridden by FDRAPPL at OPEN time, so any non-temporary name is acceptable. The name you specify will not be used by FDRAPPL, but MVS will do an exclusive ENQ on this name at job initiation so each FDRAPPL job should use unique names.
UNIT=	—specify a generic (e.g., 3490) or esoteric (e.g., CART) name to allocate the type of tape drive desired. If you have sufficient tape drives available, specifying a unit count of 2 (e.g., UNIT=(3480,2)) may reduce elapsed time (especially on 3480 cartridge drives).
VOL=	specify a volume count, e.g., VOL=(, , 255), to prevent FDRAPPL from abending if more than 5 tape volumes are required. If no volume serials are specified, FDRAPPL will call for scratch tapes; this is recommended; however, you can specify up to 255 tape volume serials if you need to select the output volumes.
LABEL=	you may want to specify RETPD= or EXPTD= to identify the expiration date of the backups (retention periods may also be specified on the DUMP statement). FDRAPPL stores this expiration date in the Application Control File entry for each data set backed up; if you have a tape management system it will honor the expiration specified. See Section 52.01 for details on retention of FDRAPPL Application Backups.
DISP=(NEW,KEEP)	is required; do not specify CATLG since FDRAPPL handles cataloging of output files internally.
FREE=CLOSE	do not specify FREE=CLOSE since it will cause FDRAPPL to fail when it tries to create a second file on the tape.

```
EXAMPLE: //TAPE1      DD   DSN=PAYROLL.BACK1,DISP=(NEW,KEEP),
           //          UNIT=3590-1,VOL=( , , 255 ),LABEL=EXPDT=99000
```

If multiple disk volumes are dumped to a given TAPEX DD, FDRAPPL will create multiple files on the tape (or tape aggregate if more than one tape volume is used), one file for each disk volume, using the naming convention in [Section 52.01](#). By default FDRAPPL will create as many as 255 files on a tape or aggregate before starting over with file 1 on a fresh scratch tape, but that limit can be overridden by the MAXFILE= operand, up to 4095 files. Larger MAXFILE= values can be used with high-capacity tapes such as IBM Magstar and StorageTek Redwood and 9840.

DCB parameters are not required and should be omitted. However, tape unit hardware compaction (sometimes called IDRC, available on most tape cartridge drives) can be requested by adding DCB=TRTCH=COMP to your DD statement; tape hardware compaction may be the default depending on local MVS options. For tapes attached by ESCON or FICON channels, Innovation recommends use of tape hardware compaction instead of FDR software compression (the COMPRESS= option).

52.04 CONTINUED

TAPE1 DD STATEMENT (continued) If you are running the StorageTek ExHPDM (High Performance Data Mover) software product, you can direct FDRAPPL backups to ExHPDM with the SUBSYS= JCL operand, e.g.,

```
//TAPE1 DD SUBSYS=(SOV,'CLASS(FDRBKUP)')
```

Please read [Section 80.33](#) and the ExHPDM program documentation for more details.

The IBM Virtual Tape Server (VTS) and similar products from other vendors are supported. In a VTS, data written to "tape" is really written to disk internal to the VTS and is later moved to high-capacity tapes such as Magstars, resulting in much better physical tape utilization. When a tape is required for input, the data is staged back to the internal disk. Don't forget that Application Backups in a VTS are primarily for on-site recovery; you may need a second copy on real tape volumes for off-site storage.

FDRAPPL output is normally directed to tape, but it is possible to direct it to disk volumes instead. Please [see Section 51.04](#) for details on the JCL required. Also [see Section 52.10](#) for notes on disk output.

WARNING: *Tapes created by FDRAPPL cannot be copied using normal copy programs. Use the INNOVATION provided program (FDRTCOPY) to copy FDRAPPL tapes.*

For SIMULATION runs, TAPE1 must specify DD DUMMY.

TAPE11 DD STATEMENT This optional DD specifies that FDRAPPL is to create a duplicate backup tape called COPY 2. This copy can be sent offsite for disaster recovery. TAPE11 does not increase the REGION requirements.

```
EXAMPLE: //TAPE11 DD DSN=PAYROLL.BACK2,DISP=(NEW,KEEP),
// UNIT=3590-1,VOL=( , , 255),LABEL=EXPDT=99000
```

will produce a tape which is a duplicate of the backups on the TAPE1 DD shown above. The copy number in the data set names created will be 2 ([See Section 52.01](#)).

SYSIN DD STATEMENT Specifies a data set containing the control statements for FDRAPPL. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

CONTINUED . . .

52.05 FDRAPPL DUMP STATEMENT

DUMP	TYPE=APPL	,ENQERR=NO
D		
	,ARCBACKUP= <u>DSF</u> NO	,ENQERR=BYPASS <u>PROCESS</u>
SIM		
	,ARCB1DSN=dsname	,HFS=QUIESCE
	,ARCB2DSN=dsname	,MAXCARDS=nnnnn
	,ARCCAT=ALL <u>NO</u> NORMAL	,MAXFILE=nnnn
	,COMPRESS=ALL COPY1 COPY2	,RETPD=dddd
		,RETPD2=dddd
	,COPY1=COPY2	
		,RTC=YES <u>NO</u>
	,DCT=YES <u>NO</u>	
		,SELTERR=NO <u>YES</u>
	,DSNENQ= <u>NONE</u> TEST USE HAVE	
		,VOLSORT= <u>YES</u> NO

There are some additional operands that can be specified on a DUMP TYPE=APPL statement, but they are rarely used with FDRAPPL. Their definitions can be found in [Section 51.05](#). They are:

CATBYPERR	DATA=	ICFCORE=
MAXBTRKS=	MAXDD=	MAXERR=
	PRINT=ABR	

For DUMP TYPE=APPL, MAXBTRKS= defaults to 65536, which disables the feature.

DUMP STATEMENT The DUMP TYPE=APPL statement invokes Application Backup. Only one DUMP statement is allowed per execution of ABR.

SIM STATEMENT If SIM is specified, FDRAPPL will perform the Application Backup in simulation mode. The TAPE1 DD statement must specify DD DUMMY, and no Application Control File will be created or written. FDRAPPL will produce a report in PRINT VTOC format showing the data sets that would be backed up. The main purpose of simulation is to check that the right data sets are being selected.

52.05 CONTINUED

OPERANDS

ARCBACKUP= **DSF** — FDRAPPL will backup the Application Control File used to record the backups for this Application Backup as the last file on the backup tapes. If the ARCB1DSN= and/or ARCB2DSN= operands are specified, those dsnames will be used for the tape file name of this backup (on TAPE1 and TAPE11 respectively). If those operands are omitted, then the backup file name is formed by taking the name of the Application Control File itself and changing the index level of ARCHIVE to ARCBKUP (for TAPE1) and ARCBKU2 (for TAPE11).

NO — FDRAPPL will not backup the Application Control File. ARCBACKUP=NO should be specified only when you are using other means to backup and recover the Application Control Files, or when the Application Backup is to be used only for onsite recovery and never at a disaster site.

Default is DSF.

ARCB1DSN=
ARCB2DSN= Specifies the data set name to be used for the backup of the Application Control File as the last file on the Application Backup tapes. This backup will be done automatically unless ARCBACKUP=NO is specified. ARCB1DSN specifies the tape file name to be used on TAPE1, and ARCB2DSN on TAPE11. These names may be any valid MVS dsnname, including a GDG relative generation, e.g.,

ARCB1DSN=PAYROLL . APPL . BACKUP (+ 1) .

These data sets will always be cataloged internally by FDRAPPL.

ARCCAT= Specifies how FDRAPPL is to handle cataloging of the tape files created by FDRAPPL.

ALL— all the files will be cataloged.

NO— no files will be cataloged, except for any backup file that exceeds 5 tape volumes and the backup of the Application Control File (see ARCBACKUP=).

NORMAL — files will be cataloged by the rules used for normal Archive Backups (it will catalog the backup file anytime the volume list changes from the previous file created, which includes the first file created, any file that requires a new output tape, and the first file created on a new output tape). However, ARCCAT=NORMAL is treated the same as ARCCAT=ALL for the TAPE1 or TAPE11 DD if EXPDT=99000 (tape management catalog control) is specified on the DD.

Default is NO.

CONTINUED . . .

52.05 CONTINUED

COMPRESS=

ALL — the backup file for both copies (TAPEx and TAPExx) is to be compressed.

COPY1 — only the backup on TAPEx DD statements will be compressed.

COPY2 — only the backup on TAPExx DD statements will be compressed.

See "Memory Requirements" in [Section 50.03](#) for the additional storage required by COMPRESS=. COMPRESS= is ignored if DCT=YES is also specified.

Default: backups will not be compressed.

COMPRESS is recommended for backups to tapes attached on parallel (bus/tag) channels. For tapes attached on ESCON or FICON channels, use of IDRC (tape hardware compression) usually results in better performance.

COPY1=COPY2

Swaps the copy numbers created on the TAPEx and TAPExx DD statements; COPY 2 will be created on the TAPEx DD and COPY 1 on the TAPExx DD (see [Section 52.01](#) for the backup data set naming conventions including the copy number). If only TAPEx is present, only COPY 2 will be created; if you send COPY 2 offsite, this can be useful when you wish to create only an offsite backup (no COPY 1 for onsite recovery).

DCT=

DCT= is valid only if you are licensed for FDR InstantBackup. It will be honored only if the disk being backed up is in a disk subsystem with the HSDM option (High Speed Data Mover). HSDM allows FDR to backup and restore the internal compressed images of disk tracks, improving backup elapsed times up to 60%. It can also be specified as DUMPCOMPRESSEDTRACK=.

YES — use HSDM for any volume where the disk hardware has the HSDM feature installed. Normal backup will be used for other volumes.

NO — do not use HSDM.

Default is NO.

Note that DCT=YES implies RTC=YES; see the description of RTC= for its benefits.

52.05 CONTINUED

DSNENQ=

Specifies whether all of the data sets selected for backup will be ENQed. See “Data Set Enqueue Option” in [Section 52.03](#) for more details.

If the ENQ fails, meaning that some other task has the data set enqueued, a warning message is issued for the data set but the data set will still be dumped unless the ENQERR=BYPASS operand is specified. A successful ENQ will prevent any other task from using the data set until the backup of that volume is complete. An ENQ failure is considered an error unless ENQERR=NO is specified, but other data sets will still be dumped. The options for DSNENQ= are:

USE — The data sets will be enqueued for the duration of the backup from this disk volume. This is the most frequently used option.

TEST — The data sets will only be tested to see if they are enqueued to another task at the time that the dump from this volume starts.

HAVE — The data sets will be enqueued for the duration of the dump. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not ENQ the data set)

RETRY (try the ENQ again)

NONE — No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: NONE.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the backup is complete. However, if the application can be sure that no other jobs will be using the data sets during the backup, use DSNENQ=NONE and avoid the ENQ overhead. You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements..

DSNENQ= should not be specified if you have multiple TAPEX DD statements, to DUMP multiple disk volumes concurrently, and multi-volume data sets may be involved.

ENQERR=

NO — If the DSNENQ= operand is used to request data set enqueues, an ENQ failure (in-use data set) will not be considered an error (see “Step Termination” in [Section 52.03](#)). Use ENQERR=NO if you want messages about active data sets but want the step to terminate normally.

Default: a DSNENQ failure will be considered an error and will cause a condition code or ABEND at step termination. This is to call attention to the error.

CONTINUED . . .

52.05 CONTINUED

ENQERR=	<p>Specifies processing if the DSNENQ= option finds that a data set is in use (enqueued):</p> <p>BYPASS – do not backup an active data set.</p> <p>PROCESS – backup a data set even if it is active (a warning message will still be produced).</p> <p>Default: PROCESS.</p> <p>NOTE: both ENQERR=NO and ENQERR=BYPASS/PROCESS may be specified on the same DUMP statement.</p>
HFS=	<p>QUIESCE invokes special processing when HFS (Hierarchical File System) data sets are backed up. HFS=QUIESCE implies DSNENQ=USE so it will first attempt to get a SYSDSN ENQ on the HFS file. If the ENQ fails, it probably means that the HFS file system is mounted to Unix System Services, so a "quiesce" call is issued to prevent updates to the HFS data set during the backup. If the quiesce fails and ENQERR=BYPASS was specified, the HFS data set will not be backed up. See "HFS" in Section 80.11 for details and security requirements.</p> <p>Note that HFS=QUIESCE implies DSNENQ=USE (described earlier) for all data sets being backed up, not just HFS data sets.</p> <p>Default: HFS data sets will not be quiesced unless HFSQUIESCE is set to YES in the FDR Global Option Table (see Section 90). If you use the default, you should unmount the HFS file system before the backup to be sure of getting a usable backup.</p>
MAXCARDS=	<p>Enables FDRAPPL to accept additional SELECT and EXCLUDE statements during this execution. You can specify values up to 65535, but values over 100 will increase the region required by FDRAPPL. If SELECT CATDSN= is used, an internal SELECT is generated for each data set/volume combination selected from the catalog, so MAXCARDS= may need to be specified if a large number of data sets are selected.</p> <p>Default is 100 statements.</p>
MAXFILE=	<p>Specifies the maximum number of files FDRAPPL will create on tape. May specify from 1 to 4095. When the maximum file number is exceeded, FDRAPPL will start a new tape using file sequence number 1. A MAXFILE= over 255 may be appropriate when you are outputting to high-capacity tapes such as IBM 3590 Magstar or StorageTek Redwood and 9840.</p> <p>Default is 255 unless overridden in the FDR Global Option table (See Section 90).</p>
RETPD= RETPD2=	<p>RETPD= specifies the number of days (1 to 9999) that COPY1 (TAPE_x) backups will be kept. FDRAPPL will calculate an expiration date from this value. If RETPD2= is omitted, this same expiration will apply to all COPY2 backups created in the same step.</p> <p>RETPD2= specifies the number of days (1 to 9999) that COPY2 (TAPE_{xx}) backups will be kept. FDRAPPL will calculate an expiration date from this value.</p> <p>See Section 52.01 for a discussion of backup retention.</p> <p>Default is 365 days, unless overridden by RETPD= or EXPDT= on individual TAPE DD statements.</p>

CONTINUED . . .

52.05 CONTINUED

RTC=

It can also be specified as READTRACKCCW=.

YES – use READ TRACK CCWs to read disk data tracks. RTC=YES also causes:

- up to 1 cylinder of disk data is read at a time.
- FDR buffers are moved above the 16MB line (about 2MB per concurrent backup), allowing more concurrent backups to be run in one step.
- the elapsed time of ABR backups when the backup data set is itself on disk is significantly improved.

NO – use other CCWs to read disk data tracks.

Default is NO.

SELTERR=

Specifies what will happen at step termination if FDRAPPL finds that a SELECT or EXCLUDE statement was never referenced (no data set on any input disk was selected by the statement):

NO — a condition code or ABEND is **not** to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.

YES — a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.

Default: YES unless overridden in the FDR Global Option Table ([See Section 90](#)).

VOLSORT=

YES — if you have more than one TAPEX DD statement in this ABR step, ABR will sort the volumes to be backed up by the last digit of their MVS device address, to attempt to balance channel and control unit utilization during concurrent backups. In other words, all volumes whose MVS address is xxx0 will be processed first, then xxx1, etc.

NO — volumes will be processed in the order that they are found during a scan of the system UCBs, which is usually (but not always) in MVS device address order.

Default is YES except that NO is forced if you have only one TAPEX DD statement.

The above does not apply to volumes selected by DISKxxxx DD statements which will always be processed first, in the order the DD statements appear in the JCL.

CONTINUED . . .

52.06 SELECT STATEMENT for FDRAPPL

SELECT	DSN=filter	,DSORG=(xx,xx..)
S	CATDSN=filter	,GDG
EXCLUDE		
X	ALLDSN	,VOL=vvvvvv
	,CATLIMITGDG=n	,VOLG=vvvvv
	,DATA=NONE	

There are some additional operands that can be specified on a SELECT statement, but they are rarely used with FDRAPPL. Their definitions can be found in [Section 51.06](#). They are:

CATALOG=	MCATALOG=
CATBYPERR	PRTALIAS
CRDAYS=	UPDATE
DATA=ALL	

- SELECT STATEMENT** The SELECT statement specifies the data sets to be dumped by FDRAPPL. SELECT CATDSN= will scan the system catalogs for the specified data sets and then scan the volumes found in the catalog to back up those data sets. SELECT DSN= or ALLDSN should also specify VOL= or VOLG= to specify the volumes to be searched. If other operands are specified (such as CRDAYS= or GDG), the data set will be backed up only if it meets those criteria.
- EXCLUDE STATEMENT** The EXCLUDE statement can be used to exclude certain data sets from a more encompassing SELECT. Since data sets are compared to the statements in the order specified, EXCLUDEs should usually precede SELECTs.
- ICF VSAM CLUSTERS** ICF VSAM clusters can be selected by FDRAPPL only by the base cluster name with the DSN= or CATDSN= operand. If selected, FDRAPPL will back up all of the components associated with this cluster which exist on the volumes processed. This includes alternate indexes and key range components. The component names will be reported by FDRAPPL followed by the base cluster name. FDRAPPL will dump to the backup file all of the associated VVR records found in the 'SYS1.VVDS' data set. For further information, [See Section 80.13](#) "VSAM Special Considerations".

OPERANDS

DSN= Specifies a fully-qualified data set name or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the VTOCs of selected volumes. VOL= or VOLG= should also be specified to identify the volumes to be scanned for these data sets.

EXAMPLES : DSN=INVENTORY.WIDGET.WAREHOUS
 DSN=CLAIMS.**
 DSN=PROD++.**.LIB*

DSN= does not have any special support for selecting GDGs. However, if it is used in conjunction with the GDG operand then only GDGs will be selected.

52.06 CONTINUED

CATDSN=

Specifies a fully-qualified data set name or a filter to be used for generic data set selection from system catalogs, as described in [Section 80.14](#).

If a fully-qualified name is specified, that name will be located in the system catalogs, and the volume serial(s) from the catalog become an implied VOL= parameter. Specification of a relative generation number for GDG data sets is supported (e.g., CATDSN=A.B(—1)).

If a filter is specified, then catalogs will be scanned for all cataloged data sets matching the filter, and they will be processed as if a SELECT CATDSN=dsname was present for each of them. It may be necessary to specify MAXCARDS=nnnnn on the DUMP statement if a large number of data sets are selected by the filter.

Additional considerations for CATDSN=filter are explained in [Section 80.14](#).

CATDSN= is supported only on SELECT statements (not on EXCLUDE).

If the VOL/VOLG= operand is also specified on a SELECT statement with CATDSN=, then only data sets cataloged to those volumes will be selected.

EXAMPLES: CATDSN=INVENTORY.WIDGET.WAREHOUS
 CATDSN=CLAIMS.**(0),VOLG=INS
 CATDSN=PROD++.**.LIB*

Normally CATDSN= will not display the data sets it selects from the catalogs; you will see the names only when FDRAPPL actually finds and selects the data sets in the VTOCs of the volumes to which they are cataloged. To display all of the data sets selected specify PCATDSN=filter.

ALLDSN

Specifies that all the data sets encountered on the volumes specified by VOL= or VOLG= are to be SELECTed or EXCLUDEd. DSN=** is equivalent to ALLDSN.

NOTE: DSN=, CATDSN= and ALLDSN are mutually exclusive. One and only one of them must be specified on each SELECT/EXCLUDE statement.

CATLIMITGDG=

May be used with CATDSN=filter to limit the selection of GDGs from the catalogs. It will not affect the selection of cataloged non-GDG data sets, but if the filter selects a GDG then:

n will cause only the most recently created "n" generations to be selected.

—n will cause only generation (—n) to be selected.

Default is that all the generations of selected GDGs will be selected unless a relative generation number is specified at the end of the filter, e.g., CATDSN=filter(-2).

DATA=

NONE — specifies that FDRAPPL will dump allocation information about the selected data sets, but will not actually back up any data tracks. Can be used when the data sets are to be allocated by FDRAPPL but the contents of the data sets will be recreated by another means. DATA=NONE must also be specified when restoring data sets that were dumped with DATA=NONE.

52.06 CONTINUED

DSORG= Specifies that the data sets is not to be selected unless its DSORG matches one of the DSORGs specified. If more than one DSORG is specified, they must be enclosed in parentheses.

VALID DSORGS are:

DA -- BDAM PS -- SEQUENTIAL AM -- ALL VSAM EF -- ICF VSAM
IS -- ISAM PO -- PARTITIONED UN -- UNDEFINED UM -- UNMOVABLE

GDG Only data sets which are generations of a Generation Data Group (GDG) will be selected.

VOL= Specifies a volume serial number (up to 6 characters) to which this statement applies. When used with CATDSN=, only data sets cataloged to this volume will be selected. If DSN= is specified, either VOL= or VOLG= is required, and those volumes are scanned for the specified data sets.

VOLG= Specifies a volume serial number prefix (up to 5 characters) to which this statement applies. When used with CATDSN=, only data sets cataloged to these volumes will be selected. If DSN= is specified, either VOL= or VOLG= is required, and those volumes are scanned for the specified data sets.

52.07 FDRAPPL RESTORE JOB CONTROL REQUIREMENTS

The following Job Control Statements are necessary to perform RESTORE from Application Backup.

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRABR resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRABR), region requirement (REGION=), and optional PARM= operand. The minimum region required is 512K. However, some restore options, especially logical restore, may increase the region requirement, so a value of 1M or 2M is recommended; REGION=0M can be specified to get the largest below-the-line region available.

If a PARM field is specified, ABR will use data specified as the first control statement, which must be a valid RESTORE statement; if the PARM data contains a slash (/), the data after the slash will be used as the second control statement (usually a SELECT). For example,

```
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=APPL,RECAT'
//FDR EXEC PGM=FDRABR,PARM='RESTORE TYPE=APPL/ SELECT DSN=A.B.C'
```

If FDRABR is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

STEPCAT or JOBCAT DD STATEMENT For non-VSAM data sets, when FDR must catalog non-VSAM data sets, they will be cataloged in that user catalog instead of in the system catalog with the matching alias. This might be useful when you are creating a test system, to catalog test copies of production data sets in a test catalog.

However, there is one important exception: if the data set being cataloged is a GDG generation, the STEPCAT/JOBCAT must contain a GDG base for that GDG; if not, it will ignore the STEPCAT/JOBCAT and catalog into the regular aliased catalog, possibly deleting other valid generations.

For VSAM clusters, the target catalog is controlled by the ICFCAT= operand, described in [Section 51.08](#). With the proper ICFCAT= option, the STEPCAT/JOBCAT may be honored.

STEPCAT/JOBCAT should not be used if any data sets being restored are SMS-managed.

SYSPRINT DD STATEMENT Specifies the output message data set; it is required. It is usually a SYSOUT data set but if it is assigned to a data set on tape or disk, this DD must specify DISP=MOD. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

ARCHIVE DD STATEMENT Specifies the Application Control File (ACF) in which all the FDRAPPL data sets to be restored were recorded.

If you are restoring at a disaster/recovery site, you may need to restore this Control File before you can restore the data sets. If the Control File was dumped to the Application Backup tape by the ARCBACKUP=DSF option, example JCL to restore it is shown in [Sections 52.11](#) and [52.12](#). You might also be recovering it by other means (such as ABR full-volume recovery).

52.07 CONTINUED

TAPE1 DD STATEMENT Optionally used to allocate a tape drive on which to mount the backup tapes from which data sets will be restored. It must specify the proper type of tape drive on which all of those backups can be read. Details are in [Section 51.07](#). However, the DYNTAPE option is recommended for Application Restores.

DYNTAPE NOTE: If DYNTAPE is specified on the RESTORE statement, this DD statement is not used and can be omitted. FDRAPPL will dynamically allocate a TAPE# DD statement for the backup device. DYNTAPE should be used if the backup is on disk, in an automated tape library (ATL or silo) or if a mixture of tape device types must be read.

SIMULATION: If SIMREST is coded, this DD usually specifies DUMMY.

SYSIN DD STATEMENT Specifies a data set containing the control statements for FDRAPPL. Usually a DD * data set. It is required, but if control statements were provided on the EXEC statement by PARM=, it can be DUMMY.

52.08 FDRAPPL RESTORE STATEMENT

RESTORE	TYPE=APPL	,MAXCARDS=nnnn
	,BLKF=nn	,NOCAT
R		,RECAT
	,COPY=n	
SIMREST		,PRESTAGE
	,DSNENQ=NONE TEST <u>USE</u> HAVE	,SELTERR=NO <u>YES</u>
	,DYNTAPE	
	,DYNTAPE2	,MSGDG=DEFERRED ACTIVE
		ROLLEDOFF <u>INPUT</u>
	,ICFCAT=ORIGINAL	
	STEP CAT ALIAS	,VRECAT

There are some additional operands that can be specified on a RESTORE statement, but they are rarely used with FDRAPPL. Their definitions can be found in [Section 51.08](#).

They are:

%FREE=	ALLOCATELIST=	BYPASSACS
BYPASSSMS	CATIFALLOC	DATA=
OPERATOR	RLSE	

**RESTORE
STATEMENT**

The RESTORE TYPE=APPL statement invokes a restore from an Application Backup. Only one RESTORE statement is allowed per execution of FDRAPPL, but any or all of the data sets for the application can be restored in one FDRAPPL step. The ARCHIVE DD statement points to an Application Control File which contains the records of data sets that were dumped with Application Backup.

**SIMREST
STATEMENT**

If SIMREST is specified, FDRAPPL will print the data set names that will be selected and the tape volumes necessary to do the restore. A restore operation is not done. SIMREST can be used to pre-pull the required tapes, or to verify that the expected data sets will be restored.

**APPLICATION
RESTORE
PROCEDURE**

If the data sets being restored already exist on the selected target disk, the restore will simply overlay them with the data from the backup tape (unless PRESTAGE is specified), so there is no need to delete the data sets before restoring them. If you are restoring at your home site, the restore will simply refresh the contents of the data sets. If they do not exist, the restore will allocate and catalog them before restoring them.

The rules for selection of a target volume for each data set are detailed in [Section 51.08](#), but briefly these tests are done in this order:

- * An NVOL= operand on the SELECT statement specifies one or more target volumes.
- * If the output data set is cataloged, FDRAPPL assumes it must exist there and uses that volume as the target.
- * The original volume serial of the data set will be used.
- * The installation may specify a global RESTORE ALLOCATION LIST which provides target volume serials based on data set name and/or original volume serial.
- * If the data set is not on the first target volume, SMS may override FDRAPPL's choice of target volume.

52.08 CONTINUED**RESTORE ICF
VSAM FILES**

FDRAPPL will restore ICF VSAM files using the base cluster name. FDRAPPL will restore each individual component associated with this cluster name, and will allocate ICF VSAM files if they do not currently exist. FDRAPPL will update the appropriate fields within the VVR for each component. ICF VSAM files except the VVDS itself are movable; except for the VVDS and catalogs they can be restored to a new name or group. If NEWGROUP= or NEWINDEX= are specified the new group name will be applied to both the cluster name and all of its components. If NEWNAME= is specified for a cluster which is not allocated, FDRAPPL will let VSAM determine names for the components. Some information contained solely in the catalog, including protection (RACF or password) and expiration date will not be updated. However, path names for alternate indexes (AIXs) and aliases of user catalogs will be restored as long as both backup and restore are done with V5.3 level 30 or above. [See Section 80.13](#) "VSAM Special Considerations" for a more detailed explanation.

OPERANDS**BLKF=**

Changes the blocksize (BLKSIZE) of certain fixed (RECFM=FB) and variable (RECFM=VB) data sets. BLKF= specifies the number of blocks per track (1 to 10) and is used to calculate a new blocksize based on the output track size (BLKF=1 results in BLKSIZE=32760). For RECFM=FB data sets the blocksize will be rounded down to a multiple of the logical record length (LRECL).

DSORG=PS (sequential) data sets will be physically reblocked during the restore; the new blocksize must be larger than the original blocksize except when restoring from 3390 to 3380.

DSORG=PO (partitioned) data sets will not be reblocked but the new blocksize will be used for new members; it must always be larger than the original blocksize.

Other data set types, including UN (undefined), DA (direct access), PSE (enhanced format sequential), POE (PDSE), HFS and all types of VSAM, will not be reblocked.

Default: data sets are not reblocked during restore; all original blocks will be restored without change, although they may be written to new locations. BLKF= is usually used when restoring to an unlike device type (e.g., 3380 to 3390) to improve utilization of the new track size but can also be used during like device restores.

COPY=

Specifies the copy of the backup from which the restore is to be attempted. "n" can be 1 or 2. COPY=2 can be specified if a duplicate tape copy (TAPExx) was created during the backup.

Default is COPY=1 unless COPY1 has expired, then COPY2 is the default if it exists. The default can be made COPY=2 in the FDR Global Option table (see [Sections 90](#)); this is useful at a disaster site where the COPY 2 backup is the only one available.

52.08 CONTINUED

DSNENQ=

Specifies whether all of the data sets being restored will be ENQed. See "Data Set Enqueue Option" in [Section 52.03](#) for more details.

If you are restoring over an existing data set and the ENQ fails, the restore will be bypassed with an error message. If the restore must allocate the output data set and the ENQ fails, no error message is issued and the restore is still done. A successful ENQ will prevent any other task from using the data set until the restore from the current backup data set is complete. An ENQ failure is considered an error but it will not prevent other data sets from being restored. The options for DSNENQ= are:

USE — The data sets will be enqueued for the duration of the restore from the current backup data set. This is the most frequently used option.

TEST — The data sets will only be tested to see if they are enqueued to another task at the time the restore starts. The data set will not be enqueued and other tasks may enqueue it while the restore is preceding.

HAVE — The data sets will be enqueued for the duration of the restore. If not available, a message (FDRW27) is issued to the MVS operator, who can respond:

WAIT (wait for the data set to become available)

NOWAIT (do not enqueue the data set)

RETRY (try the enqueue again; may result in the FDRW27 message again)

NONE— No data set ENQ will be issued.

CAUTION: This option should not be used on shared DASD unless a cross-system enqueue facility such as GRS or MIM is available and the SYSDSN QNAME is broadcast across systems. Without this capability, FDR can only determine what data sets are active on the system FDR is running on.

Default: USE. Note that NONE or TEST may allow other jobs to attempt to read the data set being restored before FDRAPPL has restored all of the data tracks.

Recommendation: use DSNENQ=USE or HAVE if you want to be sure that no other task uses the data set until the restore is complete. However, use DSNENQ=NONE when another data set by the same name on another volume may be in use (e.g., restoring data sets to an alternate SYSRES volume). You may suppress ENQs for specific data sets by the DSNENQ=NONE operand on SELECT statements.

**DYNTAPE
DYNTAPE2**

Specifies that FDRAPPL is to dynamically allocate the backup data sets using a DDNAME of TAPE#. This option should be used if the backup is on disk, in an automated tape library (ATL or silo) or a mix of backup device types will be read, such as 3480 and 3490E.

DYNTAPE2 will allocate 2 drives which will improve performance when restoring from multi-volume tape backups. DYNTAPE2 should not be used if backup files on disk might be involved.

Default: the first TAPEX DD statement found in the step JCL is used to mount all backups. The type of device assigned must be the correct type for reading all required backups.

CONTINUED . . .

52.08 CONTINUED

ICFCAT=

Applies to ICF VSAM files only. Specifies the source of the catalog name to be used if an output ICF VSAM cluster must be allocated.

ORIGINAL — use the catalog in which the original dumped cluster was cataloged. When restoring a cluster to a new name, ICFCAT=ORIGINAL is treated like ICFCAT=ALIAS, described below. If you need to catalog the output cluster into the same catalog as the input cluster but that catalog is not the one aliased for the new name, you must specify ICFCAT=STEPSCAT and supply a STEPSCAT DD statement pointing to that catalog.

STEPSCAT — use the STEPSCAT as the target catalog. If a STEPSCAT DD statement is not supplied, it will use the master catalog or the catalog which is aliased for this data set in the master catalog.

ALIAS— determine the catalog from the alias name in the master catalog. If no alias is found and the cluster is being restored to the same name, use the input cluster's original catalog. If no alias is found, and the cluster is being restored to a new name, it will use the STEPSCAT (if present in the JCL) or the master catalog. Multi-level alias (MLA) is supported.

Default: ORIGINAL, except that if the cluster is being restored to a newname (NEWGROUP or NEWINDEX specified) the default is ALIAS. If the output cluster is SMS-managed, ALIAS is forced.

ICFCAT might be needed if an Application Backup is restored on an existing system with a different catalog structure.

MAXCARDS=

Accept additional SELECT and EXCLUDE statements (over 100).

Default is 100 statements.

**NOCAT
RECAT**

NOCAT specifies that output data sets will not be cataloged. This option is ignored for ICF VSAM clusters and SMS-managed data sets, since these must always be cataloged.

RECAT specifies that non-VSAM output data sets will be cataloged even if they are currently cataloged to another volume. If a data set by that name actually exists on the volume to which it is currently cataloged, and it is SMS-managed, it will be deleted; otherwise, it will become an uncataloged data set.

Default: catalog output non-VSAM data sets only if they are not currently cataloged.

NOTE: Allocation of SMS-managed data sets will fail if they cannot be cataloged. If an SMS data set is being restored and it is currently cataloged to another volume you can either specify RECAT or delete the data set before restore.

NOCAT and RECAT are mutually exclusive. The restore will normally attempt to catalog only output data sets which it allocates (not pre-allocated) unless the CATIFALLOC operand is also specified.

RECAT is recommended for Application Restore so that the catalog is updated if they must be restored to a new volume.

52.08 CONTINUED

PRESTAGE

Output data sets which already exist on the target output volume will not be restored. This may be used to avoid restoring data sets which have already been restored. If the output data sets do not exist on the target volume, they will be allocated and restored.

Default: pre-allocated data sets will have their contents overlaid.

SELTERR=

Specifies what will happen at step termination if one or more of the SELECT or EXCLUDE statements was never referenced (no data set was selected from the Application Control File by the statement):

NO — a condition code or ABEND is not to be issued at step termination. You might use SELTERR=NO when you expect some unmatched SELECT/EXCLUDE statements, perhaps because some data sets may not exist.

YES — a condition code or ABEND will be issued at step termination to call attention to a possible control statement error.

Default: YES unless overridden in the FDR Global Option Table ([See Section 90](#)).

SMSGDG=

Specifies the status of SMS-managed GDG (Generation Data Group) data sets, if allocated by the restore.

DEFERRED, ACTIVE, or ROLLEDOFF will set the GDG to that status.

INPUT will set the GDG to the original SMS status of the GDG generation, as recorded on the backup tape. If the original GDG was non-SMS, it will be set ACTIVE if that generation is **currently** cataloged, otherwise DEFERRED.

If a GDG is restored as DEFERRED, you may need to execute an IDCAMS "ALTER ROLLIN" to make the generation active.

Default: INPUT. If you want to change this default, please contact Innovation for assistance.

CONTINUED . . .

52.08 CONTINUED

VRECAT

Allows ICF VSAM clusters to be allocated and cataloged even if they already exist in the target ICF catalog. If an attempt to define a VSAM cluster fails with a code indicating the cluster or component name already exists in the catalog, this indicates that either the cluster currently exists on another volume or the cluster is cataloged but is not on the cataloged volume. With VRECAT, the cataloged cluster will be scratched (by DELETE or, if that fails, DELETE NOSCRATCH). The define will then be re-issued. VRECAT is useful when restoring a cluster when its catalog has been restored, but the cluster on disk has not, or when restoring a cluster to a new volume.

VRECAT is ignored when:

- restoring an ICF catalog
- the restore does not include the base data component (such as restoring an alternate index on a volume by itself or a volume containing only a base index component)
- components of the cluster **do exist** on the volume to which FDR is restoring. In this case, FDR will attempt to restore on top of those existing components and VRECAT is not involved

Default: ICF VSAM clusters cannot be allocated if the cluster name already exists in the catalog (even if the catalog points to the output volume).

WARNING: VRECAT will DELETE the original cluster, with all its components, alternate indexes and PATHs, from the catalog and disks. If the DELETE fails for some reason, the DELETE NOSCRATCH may leave uncataloged components on disk.

VRECAT is recommended for Application Restore so that VSAM clusters can be restored to a new volume.

52.09 APPLICATION RESTORE SELECT STATEMENT

SELECT	DSN=filter	,NVOL=(vvvvvv,vvvvvv,...)
S	ALLDSN	,OLDBACKUP=mm
EXCLUDE	,ADATE=yyddd yyyddd/hhmmss	,PRESTAGE
X	,BLKF=nn	,VOL=vvvvvv
	,COPY=n	
	,DATA=NONE	

There are some additional operands that can be specified on a SELECT statement, but they are rarely used with FDRAPPL. Their definitions can be found in [Section 51.09](#). They are:

%FREE=	CYL=	DATA=ALL
DATACLAS=	DSNENQ=	MGMTCLAS=
NEWNAME=	NEWGROUP=	NEWINDEX=
NOCAT	NULLDATACLAS	NULLMGMTCLAS
NULLSTORCLAS	RECAT	RLSE
STORCLAS=	TRK=	VRECAT

**SELECT/
EXCLUDE
STATEMENTS**

These control statements select the data sets to be restored from Application backups. The SELECT statement identifies the individual data set name or group of data sets to be processed. The EXCLUDE statement identifies data sets from within those selected by SELECT statements which are not to be processed. ABR will select the backups of the selected data sets using information in the Application Control File, locate the indicated backup on tape, and restore it.

EXCLUDE statements should only contain the operands DSN=, or ALLDSN, and optionally VOL=.

The control statements are always scanned in the order in which they were input, so in general, EXCLUDE statements should precede SELECT statements. Since FDRAPPL will only restore data sets which are selected, EXCLUDE statements are required only to exclude certain data sets from within a larger group on a SELECT statement.

Example 1. Select all data sets with a first index of "A" except those with a second index of "B":

```
EXCLUDE DSN=A . B . **
SELECT  DSN=A . **
```

Example 2: Select all data sets backed up from volume TSO001 except those beginning with "ABC":

```
EXCLUDE DSN=ABC**
SELECT  ALLDSN , VOL=TSO001
```

NOTE: If duplicate data set names are found on the Application Control File, FDRAPPL will only select the most recently dumped data set unless ADATE=, OLDBACKUP= and/or VOL= is specified.

FDRAPPL usually restores all the data sets recorded for the application in its Application Control File using SELECT ALLDSN but selective restores can be done if required.

52.09 CONTINUED

OPERANDS

DSN= Specifies a fully-qualified data set or a filter to be used for generic data set selection, as described in [Section 80.14](#). This name or filter will be used when scanning the Application Control File.

EXAMPLES : INVENTORY.WIDGET.WAREHOUS
 DSN=CLAIMS.**
 DSN=PROD++.**.LIB*

Note: the CATDSN= operand will also be accepted on a SELECT for RESTORE, but it will be treated exactly like DSN=. It will not search system catalogs, nor will it accept GDG relative generation numbers.

ALLDSN Requests that all data sets in the Application Control File should be restored. If you are using a permanent Application Control File, this will restore the most recent copy of every data set recorded unless ADATE=, OLDBACKUP=, or VOL= is also specified.

ADATE= Only data sets which were backed up on the julian date specified (yyddd or yyyyddd) will be selected. Used to qualify a data set on the Application Control File if the same data set name was archived multiple times. 2-digit year values (yy) less than 70 are assumed to be in the 21st century, e.g., ADATE=02123 means 2002.123.

If the data set was dumped multiple times on the same day, OLDBACKUP= may also be specified to indicate which version to restore.

For Application Backups, the control file contains a time of day as well as a date. You can select by date and time, e.g., ADATE=97123/140522 or ADATE=1997123/140522.

By default, the most recent copy of each data set found in the Application Control File will be restored.

BLKF= Selected PS and PO data sets are to be reblocked during the restore; see BLKF= in [Section 52.08](#) for details.

Default: data sets are not reblocked unless BLKF= was specified on the RESTORE statement. The restore will fail if the input data set has blocks larger than the track size of the output disk.

COPY= Specifies the copy of the backup from which the restore is to be attempted. COPY=2 can only be specified if a duplicate backup (COPY 2) was created during archival or by FDRTCOPY.

Default is COPY=1 unless COPY1 was not created or has expired, then COPY2 is the default if it exists. The default can be made COPY2 in the FDR Global Option Table ([See Section 90](#)) or by COPY=2 on the RESTORE statement. This can be useful at a disaster site where COPY 2 is the only copy available.

DATA= **NONE** — the data sets selected by this SELECT statement will be allocated, but none of their data tracks will be restored. Can be used when the contents of the data sets will be recreated by another means. DATA=NONE must be specified for data sets that were dumped with DATA=NONE.

52.09 CONTINUED

NVOL=

Specifies the volume serial(s) of output disk volumes to which data sets selected by this statement are to be restored. You may specify:

- 1) A single specific volume serial, e.g.,
NVOL=ABC123
- 2) A list of specific volume serials, enclosed in parentheses, e.g.,
NVOL=(TSO001,TSO002,TSO003)
- 3) A group of volumes by placing an asterisk at the end of the volser prefix, e.g.,
NVOL=TSO*
- 4) A combination of specific and group, e.g.,
NVOL=(TSO*,PROD*,ABC001)
- 5) All online disk volumes may be selected by
NVOL=*

A list of online target volumes matching your specification is generated by scanning all disk UCBs in the system UCB chains; there is no guarantee of the order in which UCBs are found, so you cannot predict the order of the volume serials in the list. If you specify volume serials or groups which are not online, they are ignored and no error message will result.

However, if the first or only specification is a specific volume serial, it will be chosen as the first target volume, with other volumes placed after it in UCB chain order.

Also, if you are restoring a multi-volume data set (non-VSAM or SMS-managed VSAM), the volume sequence number of the piece of the data set being allocated will be used to select a specification from your list. For example, if NVOL=(A,B,C), the second piece of the data set will go to volume B. If that specification is a group, the first volume in the UCB chain matching that group will be tried. If the allocation is unsuccessful (such as insufficient free space), then other volumes in the NVOL list will be tried as described above.

The first target volume is checked to see if an output data set already exists there. If so, it restores over the existing allocation (unless PRESTAGE was specified). If not, it attempts to allocate the output data set on that volume. If the allocation fails, it will be retried on successive volumes in the list until it succeeds or until it fails on 64 volumes. If the list contains several disk device types, "like" volumes (same type as the data set being restored) will be tried first, then unlike devices.

For multi-volume data sets, a target volume is bypassed if a piece of the data set already exists there but is not the right piece, so that it will not attempt to restore the third volume of a data set on top of the first volume. When it finds a target volume in the list that does not contain a piece of the data set, it will be allocated.

Specifying multiple volsers or a volume group allows you to restore data sets in one pass even when no one volume has available space to contain them all; they will be spread across many of the target volumes.

Default: the output volume will be selected by rules defined in [Section 52.08](#). Note that when NVOL= is specified, and data sets are selected which are currently allocated and cataloged, FDRAPPL will restore them to the new volumes, and not to the volume on which they are cataloged.

On a system with SMS active, NVOL= may be ignored if the data set does not exist on the volume specified and the data set is SMS-managed.

Note: if an allocation is attempted on several volumes from your NVOL list but it fails on all of them, the message printed will usually show the allocation failure codes from the **first** volume only; failure codes from other volumes are not displayed and may be different.

CONTINUED . . .

52.09 CONTINUED

- OLDBACKUP=** If you are using a permanent Application Control File, and data sets have been backed up and recorded there more than once, OLDBACKUP= specifies which version to restore. 0 restores the most recently dumped version, 1 the next most recent, etc., up to a maximum of 127. If ADATE= is also specified, OLDBACKUP= selects from multiple backups of the data set on that date only (0 selects the last one dumped on that date, etc.).
- Default is 0.
- PRESTAGE** Specifies that selected data sets will not be restored if the output data set already exists on the first target output volume. This may be used to avoid restoring data sets which have already been restored.
- Default: restore pre-allocated data sets, overlaying the existing contents of those data sets, unless PRESTAGE was specified on the RESTORE statement.
- VOL=** Specifies the disk volume serial number from which the data set name was backed up. This operand is used to further select a data set on the Application Control File if the data set name was backed up multiple times.

52.10 MANAGING THE APPLICATION CONTROL FILE

The key to FDRAPPL is the Application Control File (ACF). When you execute FDRAPPL, each data set that is selected will be recorded in the ACF. The records in the ACF are required to automate the restore of data sets from Application Backup. The format of the Application Control File is the same as the Archive Control File described in [Section 51.01](#).

FDRAPPL creates one backup file for every disk volume processed; all selected data sets on that volume go into one backup file. Since many backup files may be produced in an FDRAPPL step if data sets are selected from more than one disk volume, the Application Control File is required so that Application Restore can find the backup file containing each requested data set.

Each ACF record consists of:

- the disk data set name
- some basic information about the data set, such as size and type
- most important, the location of the backup file (backup data set name, tape volume serials, and file sequence number). If two backups were created (COPY1 and COPY 2) both are recorded in the same record.
- expiration dates associated with each backup copy, assigned when the backup is created (See [Section 52.01](#) for details).

Whenever a restore is requested, the Application Control File which contains the records of the backup must be available. This includes off-site (disaster/recovery) restores. So your plan for using FDRAPPL must include a scheme for managing the ACF and insuring that it is available when needed.

There are two techniques that can be used with the Application Control File:

- 1) A permanent ACF for each application. Records for each backup are appended to the data already in the ACF, so that it becomes a cumulative record of Application Backups done for this application. Backups are assigned an expiration date, usually calculated from a retention period (RETPD). Records for expired backups can be deleted by the FDRARCH utility. A restore will select the most recent backup of each data set, unless you request restore of older backups. This is described in [Section 52.11](#).
- 2) A new ACF for each execution of FDRAPPL. Each ACF (usually a GDG) will contain only records of data sets backed up in that execution. Expired backups are deleted when their ACF is uncataloged (exceeds the number of generations in the GDG), so you can easily specify or change the number of backups to be kept. This is described in [Section 52.12](#).

By default, FDRAPPL will backup the Application Control File as the last file on your backup tape. The ACF in that backup will contain all of the records required to restore the application data sets from the backup just created, making that tape self-contained.

At a disaster site, you may need to restore the ACF from the backup tape before doing any restores, unless it has already been recovered by some other means (such as ABR Volume Backups).

For on-site recovery, the current ACF will usually be on disk already, so you can simply point to it and restore the data sets. However, if that ACF on disk has been lost (such as by a hardware failure), you may still need to restore it.

Note: although tape output is usually used with FDRAPPL, and all of the examples that follow reflect this. But it is possible to direct FDRAPPL to disk. When the output is on disk, the ACF is not backed up. You can use the FDRTSEL utility ([Section 60](#)) to later move those disk backups to tape; it will update the Application Control File with the new locations of the backups and it has an option (ARCBACKUP=DSF) to backup the ACF as the last file on the output tape, just like ABR.

52.11 USING A PERMANENT APPLICATION CONTROL FILE

This section describes use of FDRAPPL with a permanent Application Control File. Example jobsteps are provided. These examples are also found in the JCL library loaded during FDRAPPL's installation (see [Section 90](#)).

There will be at least one unique permanent Application Control File (ACF) for each application which uses FDRAPPL. You may need to create multiple ACFs for an application if it needs to backup different data sets at different times. The Application Control File will accumulate entries which record all of the Application Backups which have been done for that application.

The high-level index of the Application Control File data set name will also be used as the high-level index of all backup files on tape created by FDRAPPL, making it easier to identify the owner of the backup tapes and easier to define tape management rules for those tapes.

If you are restoring the data sets at your home site, you can simply reference the Application Control File, and Application Restore will restore the latest copy of every data set recorded in it.

If you are restoring at a disaster site, you must first restore the latest copy of the Application Control File. You have two options for doing this:

- You can back up and restore the Application Control File independently of the Application Backup process. One way to do this is to back up the Application Control Files for all applications using ABR Volume Backups, with full-volume and incremental backups. This means that the Application Control Files must be on volumes that are backed up by ABR Volume Backups, and which are recovered using ABR volume recovery or data set restore (see [Section 50](#)) before doing any Application Restore.
- Unless you specify ARCBACKUP=NO during the Application Backup, FDRAPPL will automatically do a DSF-type backup of the Application Control File to the last file on the backup tapes, and will catalog that backup file (the backup file can be a GDG to make it easier to recover it from various levels). You can use FDRDSF to restore the Application Control File from that backup, and then use it to restore the application data sets. It is recommended that those ACF backups be generations of a GDG (Generation Data Group) so that you can track and restore from several versions.

To successfully restore at a disaster site, you must first restore all system catalogs that are involved, including the catalogs for the application data sets, and the catalogs in which the Application Backups themselves are cataloged (based on the names of the Application Control File and the ARCB1DSN/ARCB2DSN operands).

Here are the steps required to implement FDRAPPL for a given application with a permanent ACF:

- 1) Create the GDG bases for the backups of the Application Control File.
- 2) Create and format the Application Control File with the FORMAT function of FDRARCH.
- 3) Build jobsteps that execute backups for the application. You may want to convert these into cataloged procedures (procs) which can be executed by many different applications, with symbolic parameters for data set names and other values. These steps can be inserted into application jobstreams, or can be run as separate jobs in the application cycle. But it is important to note that these backups should be run at a point in processing decided by the application, not by the Data Center.
- 4) Build jobsteps that can be used to restore data sets from the backups. These may also be converted into cataloged procedures. You may need separate jobsteps to restore the Application Control File from the backup on tape before you restore the application data sets. You may need separate procedures for onsite and offsite (disaster) restores.
- 5) Periodically execute the REORG function of FDRARCH to delete obsolete entries from the ACF and compress it.

Examples of all of these steps follow. In all the sample JCL shown, you may need to modify data set names, unit names, and other parameters to meet local requirements.

52.11 CONTINUED

**CREATE GDG
BASES IN
CATALOG**

For each application, you must execute this IDCAMS step **one time** to create the GDG bases used by that application. Every application using FDRAPPL must have a unique set of GDG bases for each Application Control File it uses. This example creates the GDG bases required by the following examples. The GDGs for the ACF backups on tape have a limit of 5, to keep 5 generations of the backup of the ACF. You can change these values to suit the application's backup and recovery requirements.

```
//BUILDGDG EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
        DEFINE GDG(NAME(PAYROLL.APPL.ACFBKP1) LIMIT(5))
        DEFINE GDG(NAME(PAYROLL.APPL.ACFBKP2) LIMIT(5))
```

**INITIALIZE THE
APPLICATION
CONTROL FILE**

To use a permanent Application Control File, you must initialize it with the FDRARCH utility, described in [Section 51.50](#). The following jobstream must be executed **one time** for each Application Control File.

The USERINDEX=YES operand is required for FDRAPPL and instructs FDRAPPL to name the backup files it creates on tape using the high-level index of the Application Control File; the security userid of the application must be authorized to create these data sets. This allows those backup files to be easily identified as belonging to this application. The rest of the name of the Application Control File is your choice.

If you use the automatic backup of the Application Control File at the end of each FDRAPPL, and you do not specify the file names for that backup with the ARCB1DSN= and ARCB2DSN= operands, the Application Control File name must have an index level of "ARCHIVE" somewhere in the name. The NODSNCK operand shown below prevents FDRARCH from insisting on an index level of "ARCHIVE".

The value for RECS= will depend on the total number of data sets that will ever be recorded in the Application Control File, considering the number of data sets in the application and the number of backups you intend to keep. However, the Application Control File can contain several hundred data sets per track (292 on a 3380, 348 on a 3390) so the space required is usually small. Any excess tracks beyond those required to hold the indicated number of data sets is released.

One record in the ACF is required for each non-VSAM data set (multiple records are required for multi-volume data sets). VSAM clusters require one record per cluster for each volume it exists on, plus one record for each component on each of those volumes (e.g., a single volume KSDS requires 3 records). This is multiplied by the number of backups that you intend to keep in the ACF, e.g., daily backup of 140 non-VSAM data sets kept for 14 days will require 1960 records (6 tracks on a 3390). Since your backup requirements may change, and since expired backups are not immediately deleted from the ACF, you should generously over-allocate the ACF.

```
//FORMAT EXEC PGM=FDRARCH
//SYSPRINT DD SYSOUT=*
//ARCHIVE DD DSN=PAYROLL.APPL.BACKUP,UNIT=SYSDA,
//        SPACE=(TRK,(10,5),RLSE),DISP=(,CATLG)
//SYSIN DD *
        FORMAT NODSNCK,RECS=10000,USERINDEX=YES
```

CONTINUED . . .

52.11 CONTINUED**FDRAPPL WITH
ARCBACKUP**

This is an example of an FDRAPPL jobstream which includes a backup of the permanent Application Control File as the last file on the backup tape. ARCBACKUP=DSF is the default and does not need to be specified.

Since both TAPE1 and TAPE11 are present, 2 copies of the backup are created, so that one can be sent offsite. The names of the backup tape files created will start with the high-level index of the ACF, but the rest of the name is controlled by FDRAPPL, as shown in [Section 52.01](#). These files will not be cataloged. They will be retained for 14 days; the tape will be expired by your tape management system at that time.

The ARCB1DSN= and ARCB2DSN= operands are included to specify the names of the tape files which will contain the Application Control File backup; they are GDGs so that multiple copies of the backup can be tracked. These files will be cataloged.

All cataloged data sets matching the SELECTs will be backed up.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1       DD     DSN=PAYBKUP.APPL1,UNIT=3490,DISP=(,KEEP),
//              VOL=(, , 255)
//TAPE11      DD     DSN=PAYBKUP.APPL2,UNIT=3490,DISP=(,KEEP),
//              VOL=(, , 255)
//SYSIN       DD     *
                DUMP TYPE=APPL,ARCB1DSN=PAYROLL.APPL.ACFBKP1(+1),
                ARCB2DSN=PAYROLL.APPL.ACFBKP2(+1),RETPD=14
                SELECT CATDSN=PAYROLL.*
                SELECT CATDSN=HOURLY.PAY*.*
```

**FDRAPPL
WITHOUT
ARCBACKUP**

This is an example of an FDRAPPL jobstream that does not backup the permanent Application Control File. If the Application Control File will be required at a disaster site, it must be recovered by other means.

Only TAPE1 is present, so only one copy of the backup is created. The names of the backup tape files created will start with the high-level index of the Application Control file, but the rest of the name is controlled by FDRAPPL, as shown in [Section 52.01](#). These files will be cataloged (ARCCAT=ALL). They will be retained for 30 days; the tape will be expired by your tape management system at that time.

All cataloged data sets matching the SELECT will be backed up. Because of CATLIMITGDG=1, for any GDGs that match the SELECT, only the current (0) generation will be selected; the operand will not affect non-GDGs, they will all be backed up.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1       DD     DSN=PAYBKUP.APPL1,UNIT=3490,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=RETPD=30
//SYSIN       DD     *
                DUMP TYPE=APPL,ARCBACKUP=NO,ARCCAT=ALL
                SELECT CATDSN=HOURLY.PAY*.*,CATLIMITGDG=1
```

52.11 CONTINUED

**DUMP FROM
MULTIPLE
VOLUMES
CONCURREN-
TLY WITH
EXHPDM**

This is an example of using FDRAPPL with the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM is invoked by the SUBSYS= operands on the TAPEX DD statements; [see Section 80.33](#) and the ExHPDM program documentation for details on the values to provide. Because there are 3 TAPEX DD statements, FDRAPPL will select data from 3 disk volumes concurrently but these concurrent backups will be interleaved into one file on a tape managed by ExHPDM. As each disk volume completes, ABR will select another volume to backup; these additional backups will be added to the same output tape file by ExHPDM, creating one large tape file containing all of the backups. This tape file will also contain the backup of the Application Control File.

```
//DUMP          EXEC  PGM=FDRABR,REGION=0M
//SYSPRINT      DD    SYSOUT=*
//SYSUDUMP      DD    SYSOUT=*
//ARCHIVE       DD    DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1         DD    DSN=PAYROLL.APPL1,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRINT1     DD    SYSOUT=*
//TAPE2         DD    DSN=PAYROLL.APPL2,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRINT2     DD    SYSOUT=*
//TAPE3         DD    DSN=PAYROLL.APPL3,DISP=(,KEEP),
//              SUBSYS=(SOV,'CLASS(FDRBKUP)')
//SYSPRINT3     DD    SYSOUT=*
//SYSIN         DD    *
                DUMP      TYPE=APPL,ARCB1DSN=PAYROLL.ACFCBP1(+1),RETPD=14,MAXERR=1
                SELECT  CATDSN=PAYROLL.**
                SELECT  CATDSN=HOURLY.PAY*.**
```

CONTINUED...

52.11 CONTINUED**RESTORING
APPLICATION
DATA SETS**

This is the jobstep necessary to restore all of the data sets that were backed up for an application. It can be used at your home site or with slight modifications at a disaster site. It assumes that the required Application Control File is on disk; at a disaster site you will need to recover the ACF before running this restore step (see next example).

It will restore the latest backup for every data set recorded in the Application Control File. By default, it will attempt to restore every data set to its original disk volume. If it cannot be allocated there, the ABR RESTORE ALLOCATION LIST, if enabled, can specify alternate volumes ([see Section 90](#)). The data sets do not need to be deleted or reallocated; FDRAPPL will simply restore over them if they exist.

If you need to restore other than the latest copy, you may need to add ADATE= or OLDBACKUP= operands to the SELECT. Also, if the application includes GDGs ABR may try to restore all recorded generations, so you may need to add ADATE= for the most recent backup date, so that older generations are not also restored (which may cause more recent generations to be deleted). You can specify DSN= on one or more SELECTS if only certain data sets are required.

If you are restoring at a disaster site, you may need to add COPY=2 to the RESTORE or SELECT statement if COPY 2 is your offsite copy. However, your installation may modify the FDR Global Option Table to make COPY 2 the default for offsite restores.

```
//RESTAPPL    EXEC    PGM=FDRABR,REGION=2M
//SYSPRINT     DD      SYSOUT=*
//SYSUDUMP     DD      SYSOUT=*
//ARCHIVE      DD      DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SYSIN        DD      *
               RESTORE TYPE=APPL,RECAT,VRECAT,DYNTAPE
               SELECT ALLDSN
```

52.11 CONTINUED

**RESTORING
THE ACF AND
APPLICATION
DATA SETS**

If the Application Control File required is **not** already on disk, as it might be a disaster site, or after a hardware error at your home site, this two step procedure can be used to restore the ACF, and then to restore the application data sets recorded in it. The input tape is mounted only once unless the backup used multiple tape volumes.

This jobstream assumes that you have already restored the catalogs into which the backup of the Application Control File was cataloged. (If not, it is possible to restore the Application Control File by specifying DSN=, and VOL=SER= and UNIT= and LABEL=, if that information is available at the recovery site).

This example restores the COPY 2 backup of the ACF, plus the COPY 2 backups of all the application files; this is usually used at a disaster site, since COPY 2 is usually sent offsite.

Since the Application Control File is backed up in FDRDSF format, FDRDSF is used to restore it. This example will restore it to its original volume, and allocate it there if necessary. Since the Application Control File is the only data set in this backup, SELECT ALLDSN will restore it. If the original volume is not available, you may need to specify the NVOL= operand to specify a new target volume.

See the previous example for notes on the Application Restore step.

```
//RESTAF      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=PAYROLL.APPL1.ACFBKP2(0),DISP=(OLD,PASS)
//SYSIN       DD     *
      RESTORE TYPE=DSF,RECAT
      SELECT ALLDSN
//RESTAPPL    EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1       DD     DSN=PAYROLL.APPL1.ACFBKP2(0),DISP=(OLD,KEEP)
//SYSIN       DD     *
      RESTORE TYPE=APPL,RECAT,VRECAT,COPY=2
      SELECT ALLDSN
```

CONTINUED . . .

52.11 CONTINUED

**APPLICATION
CONTROL FILE
MAINTENANCE**

With permanent Application Control Files, a job must be run periodically against each Application Control File to reorganize it, primarily to eliminate expired entries. This job should usually be run once a week, but it may be run more or less frequently depending on how frequently FDRAPPL is run and the number of data sets dumped.

This example will delete all expired entries, those past their expiration date. This assumes that a retention period (RETPD=) or expiration date (EXPDT=) was specified at backup time as shown in the earlier examples. It is possible to specify other criteria for deletion such as MAXOCCURRENCES=nn to retain "nn" copies of each data set.. In particular, if the application includes GDGs, you may want to specify MAXGENERATIONS=nn, so that only the appropriate number of generations will be retained.

TAPE1 specifies a data set used to take a temporary backup of the Application Control File during the reorganization. It must be large enough to hold the entire contents of the ACF.

Details on the REORG function of FDRARCH are found in [Section 51.50](#).

```
//REORG      EXEC  PGM=FDRARCH,REGION=1M
//SYSPRINT   DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//ARCHIVE    DD    DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1      DD    UNIT=SYSDA,SPACE=(TRK,(10,5))
//SYSIN      DD    *
              REORG  DUMPDEVICE=DISK,SORT=YES,SORTALLOC=YES
```

52.12 USING A GDG APPLICATION CONTROL FILE

This section describes use of FDRAPPL with a GDG Application Control File. Example jobsteps are provided. These examples are also found in the JCL library loaded during FDR's installation (see [Section 90](#)).

In this technique, a new Application Control File is created every time the FDRAPPL jobstream is run. The ACF is created as a generation of a GDG (Generation Data Group) so that multiple versions of it can be kept. A unique GDG must be created for every FDRAPPL but the name of the GDG can be any that meets your needs.

You must define the GDG base for the ACF with a LIMIT of 1 or more, reflecting the number of versions of that ACF that you want to retain on disk. The GDG base should have the SCRATCH attribute so that old versions are scratched from disk automatically when new versions are created.

If you use the default of ARCBACKUP=DSF to make a copy of the Application Control File on tape, those backup files will usually also be a GDG. You must also create the GDG bases for those backups, with LIMITs which reflect the number of backups to be kept for the application. If you define the Application Control File GDG with LIMIT(1) so that only the latest version is kept on disk, it can be used for on-site restores from the latest backup. If you need to restore from an older backup or at a disaster site, you must restore the ACF from the backup tape before executing restores.

These GDG Application Control Files do not require any maintenance, since each contains only the records from one execution of FDRAPPL and are automatically deleted when obsolete.

The high-level index of the Application Control File data set name will also be used as the high-level index of all backup files on tape created by FDRAPPL, making it easier to identify the owner of the backup tapes and easier to define tape management rules for those tapes.

If you are restoring the data sets at your home site, you can simply reference the proper generation of Application Control File (usually "(0)" to restore the latest backup) and Application Restore will restore every data set recorded in it. If you are restoring at a disaster site, you must first restore the latest copy of the Application Control File. You may also need to do so if you are restoring from other than the current backup, or if the copies of the ACF on disk have been lost. You have two options for doing this:

- You can back up and restore the Application Control File independently of the Application Backup process. One way to do this is to back up the Application Control Files for all applications using ABR Volume Backups, with full-volume and incremental backups. This means that the Application Control Files must be on volumes that are backed up by ABR Volume Backups, and which are recovered using ABR volume recovery or data set restore (see [Section 50](#)) before doing any Application Restore.
- Unless you specify ARCBACKUP=NO during the Application Backup, FDRAPPL will automatically back up the Application Control File used as the last file on the backup tapes, and will catalog that backup file (the backup file should also be a GDG to make it easier to recover it from various levels). You can restore the Application Control File from that backup, and then use it to restore the application data sets.

To successfully restore at a disaster site, you must first restore all system catalogs that are involved, including the catalogs for the application data sets, and the catalogs in which the Application Backups themselves are cataloged (based on the names of the Application Control File and the ARCB1DSN/ARCB2DSN operands).

52.12 CONTINUED

Here are the steps required to implement FDRAPPL for a given application with a permanent ACF:

- 1) Create the GDG bases for the Application Control file and the backups of the Application Control File.
- 2) Build jobsteps that execute backups for the application. Each step will create a new generation of the ACF. You may want to convert these into cataloged procedures (procs) which can be executed by many different applications, with symbolic parameters for data set names and other values. These steps can be inserted into application jobstreams, or can be run as separate jobs in the application cycle. But it is important to note that these backups should be run at a point in processing decided by the application, not by the Data Center.
- 3) Build jobsteps that can be used to restore data sets from the backups. These may also be converted into cataloged procedures. You may need separate jobsteps to restore the Application Control File from the backup on tape before you restore the application data sets. You may need separate procedures for onsite and offsite (disaster) restores.

Examples of all of these steps follow. In all the sample JCL shown, you may need to modify data set names, unit names, and other parameters to meet local requirements.

CREATE GDG BASES IN CATALOG

For each application, you must execute this IDCAMS step one time to create the GDG bases used by that application. Every application using FDRAPPL must have a unique set of GDG bases for each Application Control File it uses. This example creates the GDG bases required by the following examples. It creates the disk Application Control File GDG with a limit of 1 generation, so that only the most recent version is kept on disk, but the tape backup GDGs have a limit of 5, to keep 5 generations of the backup of the ACF. You can change these values to suit the application's backup and recovery requirements.

```
//BUILDGDG EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
    DEFINE GDG(NAME(PAYROLL.APPL.BACKUP) LIMIT(1) SCRATCH)
    DEFINE GDG(NAME(PAYROLL.APPL.ACFBKP1) LIMIT(5))
    DEFINE GDG(NAME(PAYROLL.APPL.ACFBKP2) LIMIT(5))
```

52.12 CONTINUED

**FDRAPPL WITH
ARCBACKUP**

This is an example of an FDRAPPL jobstream which includes a backup of the Application Control File GDG as the last file on the backup tape. ARCBACKUP=DSF is the default with TYPE=APPL and does not need to be specified.

The ARCHIVE DD creates the Application Control File for this backup step on disk.

Since both TAPE1 and TAPE11 are present, 2 copies of the backup are created, so that one can be sent offsite. The names of the backup tape files created will start with the high-level index of the ACF (PAYROLL), but the rest of the name is controlled by FDRAPPL, as shown in [Section 52.01](#). These files will not be cataloged even though the JCL specifies EXPDT=99000 (catalog control for many tape management systems).

The ARCB1DSN= and ARCB2DSN= operands are included to specify the name of the tape file which will contain the Application Control File backup; they are GDGs so that multiple copies of the backup can be tracked. FDRAPPL will create and catalog these files internally, as the last file on each of the output tapes. Since they are cataloged and under catalog control, your tape management system will retain the entire tape set until the ACF backup generation is deleted (rolls out of the GDG).

All cataloged data sets matching the SELECTs will be backed up.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP(+1),DISP=(NEW,CATLG),
//              UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)      (see note)
//TAPE1       DD     DSN=PAYBKUP.APPL1,UNIT=CART,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=EXPDT=99000
//TAPE11      DD     DSN=PAYBKUP.APPL2,UNIT=CART,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=EXPDT=99000
//SYSIN       DD     *
              DUMP TYPE=APPL,ARCB1DSN=PAYROLL.APPL.ACFBKP1(+1),
              ARCB2DSN=PAYROLL.APPL.ACFBKP2(+1)
              SELECT CATDSN=PAYROLL.**
              SELECT CATDSN=HOURLY.PAY*.*
```

Note: Depending on your system and local requirements, you may need to add a DCB= keyword or define a MODEL DSCB to create a GDG generation.

CONTINUED . . .

52.12 CONTINUED

**FDRAPPL
WITHOUT
ARCBACKUP**

This is an example of an FDRAPPL jobstream that does not backup the GDG Application Control File. If the Application Control File will be required at a disaster site, it must be recovered by another means.

Only TAPE1 is present, so only one copy of the backup is created. The names of the backup tape files created will start with the high-level index of the Application Control file, but the rest of the name is controlled by FDRAPPL, as shown in [Section 52.01](#). These files will be cataloged (ARCCAT=ALL). They will be retained for 30 days; the tape will be expired by your tape management system at that time.

All cataloged data sets matching the SELECT will be backed up. Because of CATLIMITGDG=1, for any GDGs that match the SELECT, only the current (0) generation will be selected; the operand will not affect non-GDGs, they will all be backed up.

```
//BACKUP      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSPRIN1    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP(+1),DISP=(,CATLG),
//              UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)      (see note)
//TAPE1       DD     DSN=PAYBKUP.APPL,UNIT=3490,DISP=(,KEEP),
//              VOL=(, , 255),LABEL=RETPD=30
//SYSIN       DD     *
              DUMP   TYPE=APPL,ARCBACKUP=NO,ARCCAT=ALL
              SELECT CATDSN=HOURLY.PAY*.*,CATLIMITGDG=1
```

Note: Depending on your system and local requirements, you may need to add a DCB= keyword or define a MODEL DSCB to create a GDG generation.

**DUMP FROM
MULTIPLE
VOLUMES
CONCURREN-
TLY WITH
EXHPDM**

This is an example of using FDRAPPL with the ExHPDM (High Performance Data Mover) software product from StorageTek. ExHPDM is invoked by the SUBSYS= operands on the TAPEX DD statements; [see Section 80.33](#) and the ExHPDM program documentation for details on the values to provide. Because there are 3 TAPEX DD statements, FDRAPPL will select data from 3 disk volumes concurrently but these concurrent backups will be interleaved into one file on a tape managed by ExHPDM. As each disk volume completes, ABR will select another volume to backup; these additional backups will be added to the same output tape file by ExHPDM, creating one large tape file containing all of the backups. This tape file will also contain the backup of the Application Control File.

```
//DUMP        EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP(+1),DISP=(NEW,CATLG),
//              UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)      (see note above)
//TAPE1       DD     DSN=PAYBKUP.APPL1,DISP=(,KEEP),
//              LABEL=EXPDT=99000,SUBSYS=SOV
//SYSPRIN1    DD     SYSOUT=*
//TAPE2       DD     DSN=PAYBKUP.APPL2,DISP=(,KEEP),
//              LABEL=EXPDT=99000,SUBSYS=SOV
//SYSPRIN2    DD     SYSOUT=*
//TAPE3       DD     DSN=PAYBKUP.APPL2,DISP=(,KEEP),
//              LABEL=EXPDT=99000,SUBSYS=SOV
//SYSPRIN3    DD     SYSOUT=*
//SYSIN       DD     *
              DUMP   TYPE=APPL,ARCB1DSN=PAYROLL.APPL.ACFBKP1(+1),MAXERR=1
              SELECT CATDSN=PAYROLL.**
              SELECT CATDSN=HOURLY.PAY*.*
```

52.12 CONTINUED

**BACKUP ONLY
ALLOCATION**

Sometimes you need to allocate certain application data sets at a disaster site, but plan to restore their contents by another means. For example, you might plan to recover some data bases with a data base recovery utility, but that utility needs to have the data bases pre-allocated. FDRAPPL can do this with the DATA=NONE operand.

When DATA=NONE is specified on a SELECT statement, FDRAPPL will back up only the description of the data set, sufficient to allocate it, but will not back up any data tracks. When restoring the data sets, DATA=NONE must also be specified on the SELECT statement so that FDRAPPL will not attempt to restore the missing data tracks. For this reason, it is easiest to dump the DATA=NONE data sets into a separate Application Control File, as shown here.

Since this backup will normally be required at a disaster site, COPY1=COPY2 is used to create only COPY 2 using the TAPE1 DD statement; all COPY 2s, including those for regular Application Backups, can be sent offsite.

It is possible to include DATA=NONE allocation-only backups in the same backup as regular Application Backups. In this case be sure to put the SELECTs with DATA=NONE before other SELECTs, if they are a subset, and be sure to include similar SELECTs in the restore jobs.

The Application Restore jobstreams shown in this section is used to restore these allocations, except that DATA=NONE must be added to the SELECT.

```
//BACALLOC      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSPRINT1     DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=PAYROLL.ALLOC.BACKUP(+1),DISP=(,CATLG),
//              UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)      (see note)
//TAPE1         DD     DSN=PAYBKUP.ALLOC1,UNIT=3490,DISP=(,KEEP),
//              VOL=(,,255)
//SYSIN         DD     *
                DUMP TYPE=APPL,ARCB1DSN=PAYROLL.ALLOC1.BACKUP2(+1),
                RETPD=14,COPY1=COPY2
                SELECT CATDSN=PAYROLL.DB.**,DATA=NONE
```

Note: Depending on your system and local requirements, you may need to add a DCB= keyword or define a MODEL DSCB to create a GDG generation.

**RESTORING
APPLICATION
DATA SETS**

This is the jobstep necessary to restore all of the data sets that were backed up for an application. It can be used at your home site or with slight modifications at a disaster site. It assumes that the required Application Control File is on disk; at a disaster site you will need to recover the ACF before running this restore step (see next example).

It will restore the backup for every data set recorded in the most recent generation of the Application Control File. By default, it will attempt to restore every data set to its original disk volume. If it cannot be allocated there, the ABR Restore Allocation List, if enabled, can specify alternate volumes (see [Section 90](#)). The data sets do not need to be deleted or reallocated; FDRAPPL will simply restore over them if they exist.

If you are restoring at a disaster site, you may need to add COPY=2 to the RESTORE or SELECT statement if COPY 2 is your offsite copy. However, your installation may modify the FDR Global Option Table to make COPY 2 the default for offsite restores.

```
//RESTAPPL      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=PAYROLL.APPL.BACKUP(0),DISP=SHR
//SYSIN         DD     *
                RESTORE TYPE=APPL,RECAT,VRECAT,DYNTAPE
                SELECT ALLDSN
```

CONTINUED . . .

52.12 CONTINUED

**RESTORING
THE ACF AND
APPLICATION
DATA SETS**

If the Application Control File required is **not** already on disk, as it might be a disaster site, or when you need to restore from other than the most recent backup, this two step procedure can be used to restore the ACF, and then to restore the application data sets recorded in it. The input tape is mounted only once unless the backup used multiple tape volumes.

This jobstream assumes that you have already restored the catalogs into which the backup of the Application Control File was cataloged. (If not, it is possible to restore the Application Control File by specifying DSN=, and VOL=SER= and UNIT= and LABEL=, if that information is available at the recovery site).

This example restores the COPY 2 backup of the ACF, plus the COPY 2 backups of all the application files; this is usually used at a disaster site, since COPY 2 is usually sent offsite.

Since the Application Control File is backed up in FDRDSF format, FDRDSF is used to restore it. This example will restore it to its original volume, and allocate it there if necessary. Since the Application Control File is the only data set in this backup, SELECT ALLDSN will restore it. If the original volume is not available, you may need to specify the NVOL= operand to specify a new target volume.

See the previous example for notes on the Application Restore step.

```
//RESTAF      EXEC   PGM=FDRDSF,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//TAPE1       DD     DSN=PAYROLL.APPL.ACFBKP2(0),DISP=(OLD,PASS)
//SYSIN       DD     *
    RESTORE TYPE=DSF,RECAT
    SELECT DSN=PAYROLL.APPL.*,NEWNAME=PAYROLL.APPL.BACKUP(0)
//RESTAPPL    EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//ARCHIVE     DD     DSN=PAYROLL.APPL.BACKUP(0),DISP=SHR
//TAPE1       DD     DSN=PAYROLL.APPL.ACFBKP2(0),DISP=(OLD,KEEP)
//SYSIN       DD     *
    RESTORE TYPE=APPL,RECAT,VRECAT,COPY=2
    SELECT ALLDSN
```

52.20 REPORTING ON APPLICATION BACKUPS

You can report on the data sets recorded in an Application Control File using:

- program FDRABRP with its PRINT ARCHIVE function (see [Section 53](#)).
- program FDREPORT with its DATATYPE=ARCHIVE option (see [Section 54](#))
- the SRS ISPF Dialog (see [Section 54](#))

Each will display the data set names, data set information, and the location and expiration of the backup files. SRS can also be used to restore data sets from the Application Backup; see the SRS example in [Section 52.21](#).

FDRABRP The ABR report utility, FDRABRP, can produce a fixed-format report on data sets recorded in a single Application Control File. This example shows details of every data set in the ACF. See [Section 53.04](#) for additional operands that you might use to limit the scope of the report. The ARCHIVE DD statement points to the Application Control File. Since SORT=YES is the default in batch, to display information in data set name order, SORT DD statements may be required as shown, depending on your system SORT product.

```
//PRINT          EXEC  PGM=FDRABRP
//SYSPRINT       DD    SYSOUT=*
//SYSOUT         DD    SYSOUT=*
//ABRMAP         DD    SYSOUT=*
//ARCHIVE        DD    DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SORTLIB        DD    DSN=SYS1.SORTLIB,DISP=SHR
//SORTWK01       DD    SPACE=(TRK,50,,CONTIG),UNIT=DISK
//SORTWK02       DD    SPACE=(TRK,50,,CONTIG),UNIT=DISK
//SORTWK03       DD    SPACE=(TRK,50,,CONTIG),UNIT=DISK
//SYSIN          DD    *
PRINT           ARCHIVE
```

The report produced will be similar to:

-DSK VOL-	*-DATE OF--*	D/S	REC	BLOCK		LBP	BK	*-DSN	TAPE	INFO--*					
-DATA SET NAME--	SERIAL	S	TY	ARCHIVE	EXPDT	ORG	FM	SIZE	LRECL	ALLOC	R/TRK	*TY	CN	SUFFIX	FN VOL-SERIAL(S)
TEMPEST.EXTRACT.GDGOX.	SMS812	2	0E	97.189	98.189	PS	FB	23440	80	1C	9	84	1	B197189A	1 AB0003
G0002V00															
TEMPEST.MASTER	SMS812	2	0E	97.189	98.189	PS	FB	23440	80	1C	9	84	1	B197189A	1 AB0003

FDREPORT The ABR generalized reporting utility, FDREPORT, can report on data sets recorded in a single Application Control File. The format of the report and the information included in it are customizable. This example will report on the basic descriptive data of each data set (volume, dsorg, size, etc.), plus information on the backup of each data set (backup data set name, volume, file, and date of backup).

```
//PRINT          EXEC  PGM=FDREPORT,REGION=2M
//SYSPRINT       DD    SYSOUT=*
//SYSOUT         DD    SYSOUT=*
//ABRMAP         DD    SYSOUT=*
//ARCHIVE        DD    DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SYSIN          DD    *
REPORT FIELD=(DEFAULTS,BKINFO)
PRINT DATATYPE=ARCHIVE
```

The generated report will look similar to the PRINT ARCHIVE report above. Please see [Section 54](#) for information on further customizing this report.

52.20 CONTINUED

**REPORTING
ON MULTIPLE
APPLICATION
CONTROL
FILES**

The preceding examples reported only on the contents of one Application Control File at a time. You may wish to produce a combined report on the contents of multiple ACFs. For example, if you create the ACF for an application as a GDG, and you retain more than one generation on disk (LIMIT for the GDG base greater than 1), you might want to report on all backups in all of those ACFs. This special FDREPORT step will do this.

FDREPORT first identifies all the generations that exist for a specified Application Control File name, extracts the data from each into a combined reporting file, then reports on that combined data. It sorts on data set name and disk volser. RPTYPE=ARCHIVE requests the same report format as produced by the PRINT ARCHIVE function of FDRABRP (see earlier example) but you may choose to customize the report as shown in [Section 54](#).

```
//REPTAPPL EXEC PGM=FDREPORT,REGION=2M
//SYSPRINT DD SYSOUT=*
//ABRMAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUT2 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SYSPUNCH DD UNIT=SYSDA,SPACE=(TRK,(2,1))
//SYSIN DD *
    DEFAULT SORTALLOCATE=YES,ENABLE=MODOUTPUT
* FIND ALL GENERATIONS OF THE SPECIFIED APPLICATION CONTROL FILE
* GENERATE FDREPORT STATEMENTS TO PROCESS THEM.
  XSELECT XDSN=PAYROLL.APPL1.BACKUP.G*,DSORG=DA
  PUNCH FDRLIB=MASK
  PRINT RPTYPE=SELPCH,DATATYPE=CATVTOC
  CANCEL
* EXECUTE THE GENERATED FDREPORT STATEMENTS TO CREATE A COMBINED
* EXTRACT FILE
  EXECUTE FDRLIB=SYSPUNCH
  CANCEL
* GENERATE THE COMBINED REPORT FROM THE EXTRACT FILE
  SORT FIELD=(DSN,VOL)
  PRINT DATATYPE=EXTRACT,RPTYPE=ARCHIVE
//MASK DD *
  DEFAULT ARCDN=<DSN>,COPY=BOTH
  PRINT DATATYPE=ARCHIVE,RPTYPE=DATA
```

52.21 ISPF SUPPORT FOR APPLICATION BACKUPS

The ABR ISPF dialogs support data set restores for FDRAPPL.

**SRS ISPF
DIALOG**

SRS stands for "Search, Report, and Services". It is described in detail in [Section 54](#). It allows users to search for information on data sets from a variety of sources, display the information in a easy-to-use format, and optionally to execute various TSO and ABR functions against the data sets displayed.

For Application Backups, it can display backup information recorded in the Application Control File for the data sets selected. Simply by typing the RESTORE command on the proper line, users can request that a data set be restored.

From your ISPF main menu, if you enter "A.S.1", you will receive this SRS panel:

----- DATASET SELECTION: ABRBKUP ----- LINE 1/16 COL 4:6/10									
COMMAND ==>					SCROLL ==> HALF				
ENTER SELECTION CRITERIA									
	Read	Save	Submit	Find	Locate	Extract	Options	Help	
FIELD	SELECTION VALUE							REPORT	SORT

DSNAME	==> **								
VOL	==>							2	
SOURCE	==> appl (Catalog Volume Archive Appl Scratch Extract)								
ARCDN	==> payroll.appl.backup(0)								
DEVTYPE	==>							3	
DSORG	==>							4	
BKINFO	==>							6	
SIZE	==>							5	

In this example, the user has requested all data sets (**), has specified several attributes of the selected data sets to display (including FDRAPPL backup information), and had requested that all recorded backups of each data set be displayed (OLDBACKUP=ALL). Information from the specified Application Control File has been displayed.

Note that all of the SRS report and selection fields shown may not display on the screen at once, you may need to scroll up and down to find the SRS fields you need. [Section 54](#) includes information and examples on customizing SRS to display just the information required for a particular function, such as restore from Archive Backup.

In response to this request, SRS will search the Archive Control File for the data sets requested and display their backup information in a format similar to:

----- DATASET LIST: DSLIST ----- LINE 1/42 COL 3:7/13									
COMMAND ==>					SCROLL ==> HALF				
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help
COMMAND	ENTRY NAME				VOLSER	DEVTYPE	DSO	ALLOC	BKDATE
-----	-----				-----	-----	-----	-----	-----
restore	PAYROLL.MASTER				SMS812	3380	**	4	1997.210

52.21 CONTINUED

The RESTORE command will take you to another panel where you can set options for the restore:

```

----- FDRSRS - Archive Restore ----- Row 1 to 2 of 2
COMMAND ==> submit                                SCROLL ==> HALF

  Edit JCL   Submit JCL   FG - execute in the foreground

Operands for RESTORE TYPE=APPL statement (section 52.08):
==> RESTORE TYPE=APPL,RECAT,VRECAT,DT,ARCDSN=PAYROLL.APPL.BACKUP(0)

DSNAME / Filter ==> 'PAYROLL.MASTER'
Volume Serial   ==> SMS812                      Backup date ==> 97210
New DSNAME      ==>
or NEWINDEX     ==>
New Volser(s)   ==>                               Copy ==> 1
Operands for SELECT DSN= statement (section 52.09):
==> NOTIFY=BAB

```

The proper volume and backup date have already been filled in. If required, you can enter a new name or NEWINDEX for renaming the restored data set; you can also specify a new target volser if it is not to be restored to the original volume.

The user has the option to submit a batch FDRAPPL jobstream to perform the restore (EDIT or SUBMIT commands), or to attach ABR under TSO and execute the restore in the foreground (FG command). Note that foreground (FG) restores from tape will work only if the TSO user is authorized to mount tapes. Your installation may have modified this panel to limit the choices available, or it may execute a particular command (e.g., RQ) automatically without displaying this panel.

Because of the NOTIFY= operand, your TSO session will receive a message when the restore of each data set is complete.

53.01 STANDARD REPORTING FACILITY (FDRABRP)

FDRABRP is a report facility for ABR which provides simple fixed-format reports on various types of ABR information, including volume status, data sets, scratched data sets, and ABR backups. It also includes a function for reporting on the contents of backups produced by any component of the FDR system.

**FDRABRP
REPORTS**

These reports are available from FDRABRP:

Archive Report – prints the data sets that have been archived from disk to tape.

Catalog Report – prints the ABR Catalog, listing the disk volumes and their related backup tapes.

Scratch Report – prints the ABR catalog, listing the data sets that been scratched and their related backup tapes.

VTOC Report – prints the contents of the VTOC for one or more disk volumes in selected format. The entire volume or selected data sets may be printed.

Tape VTOC Report – prints the contents of the formatted VTOC at the beginning of FDR, DSF and ABR backup tapes. The data sets which were dumped to this backup can be displayed in ABR or IEHLIST format. For ICF VSAM clusters on the backups, a simulated IDCAMS LISTCAT listing can be produced.

Backup Reports – prints the backup information (backup date, volumes backup to, etc.) for the data sets selected. The most current backup and optionally previous (OLDBACKUP) information can be requested.

Data Set Utilization by ID – reports by index on the number of data sets owned by type (ex: PS, PO) and the number of tracks owned by these data sets. Also, the user can request aging statistics by index. The index can be the highest level index or some other index level within the data set name.

Wasted Space Reports – prints volume utilization reports. It will report on the percentage of free space on a volume and the amount of used space within PDS and sequential data sets.

Volume Status Report – lists the ABR status of all or selected disk volumes ONLINE at execution time.

The only control statement supported by FDRABRP is the PRINT statement, with various operands to request the reports above.



FDRABRP provides basic preformatted (canned) reports for tracking ABR usage. They are simple to generate but have limited selection capability and inflexible output format. Innovation strongly recommends that you consider using the SRS ISPF dialogs documented in [Section 54](#) for most of your reporting requirements.

The SRS dialogs allow you to report not only on data set activity but also to act on this information within the same panel. For example, if you run the FDRABRP backup report to get information on the backup history of a data set and then you decide to restore the data set you would have to enter the ISPF restore dialog or edit a batch jobstream and reenter the data set name, GEN, CYCLE, etc. With SRS you can request the backup information, select the backup version you want to restore, and simply enter the request on the right line of the display.

SRS reports can be customized to the requirements of each user.

For batch reporting, consider using FDREPORT, the ABR generalized report writer, also documented in [Section 54](#).

53.01 CONTINUED

ISPF DIALOG SUPPORT The ABR ISPF dialogs allow you to generate some of the FDRABRP reports in this section by filling in required options on a panel. ISPF option A.1 takes you to a panel where you can select the type of report desired, as shown below:

PANEL A.1

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM - FDRABR REPORT PANEL -----

REPORT OPTION ==>
BLANK   - ARCHIVE                               ENTER 'C' TO CHANGE FORMAT ==>
 2      - BACKUP                                PRINT ALL AVAILABLE BACKUPS
 3      - SCRATCH                               CREATE DETAIL AND SUMMARY REPORTS
 4      - CATALOG                              REPORT DIRECT TO TSO
 5      - VOLUME STATUS                         DISPLAY IN TSO FORMAT
 6      - FDREPORT                             CREATE 58 LINES PER PAGE

FDREPORT NAME ==>

PROJECT ==>
LIBRARY ==>
QUALIFIER =>

OTHER DATA SET NAME ==>
OTHER DATA SET GROUP ==>

VOLUME SERIAL ==>
VOLUME GROUP ==>

```

Here you select the type of report on the REPORT OPTION line.

For the ARCHIVE, BACKUP, and SCRATCH reports, you can specify the data set name in either:

- TSO format (PROJECT, LIBRARY, QUALIFIER). If QUALIFIER or LIBRARY and QUALIFIER are omitted, the remaining levels are treated as a data set name prefix (DSG=).
- other data set name or prefix (group). If you do not put the name or group in quotes, your TSO userid will be assumed as a prefix. For example, if you put JCL.CNTL for OTHER DATA SET NAME, it will actually display "userid.JCL.CNTL".

A DASD volume serial or volser prefix (group) can optionally be specified for these reports.

For the CATALOG and VOLUME STATUS (VOLSTAT) reports, a volume serial or group is normally required (if omitted, it will report on all of your online DASD which may take considerable time).

Details on these reports are found in the rest of this section.

Option 6 (FDREPORT) is a way of executing predefined FDREPORT input, but the SRS ISPF dialogs provide a much better way of doing so. FDREPORT and SRS are documented in [Section 54](#).

CONTINUED . . .

53.02 FDRABRP JCL REQUIREMENTS

The following Job Control Statements are required to execute FDRABRP:

EXEC STATEMENT	Must specify the name of the ABR print utility, FDRABRP. The EXEC statement may also contain the region requirements of 256K. You can also specify the PARM= parameter; if present, it will be processed as the first control statement.
STEPCAT DD STATEMENT	A STEPCAT may be required for the VTOC or backup report, if ICF VSAM files are present and their high level index is not aliased in the master catalog.
SYSPRINT DD STATEMENT	Specifies the primary output message data set. This is a required DD statement for all PRINT functions and is usually a SYSOUT data set.
ABRMAP DD STATEMENT	Specifies the Report data set. when ABRMAP is not found within the JCL stream, the reports will be output to the SYSPRINT data set. Usually a SYSOUT data set.
ABRSUM DD STATEMENT	Required by the PRINT VTOC report by data set index. Usually a SYSOUT data set.
VSAMPRT DD STATEMENT	Required by the PRINT TVTOC command when VSAM=YES or VSAM=DUMP is present. Usually a SYSOUT data set.
ARCHIVE DD STATEMENT	Specifies the ABR ARCHIVE control file required for the PRINT ARCHIVE statement. MVS users need not specify this DD statement if the ARCHIVE control file is the standard file name specified in the ABR option table and is cataloged. The control file will be dynamically allocated to DDNAME ARCHIVE#.
DISKxxxx DD STATEMENT(S)	<p>Optionally specifies the DASD volumes(s) to be processed. xxxx is any combination of valid alphanumeric digits (0-9, A-Z). They should look like:</p> <pre>//DISK1 DD UNIT=DISK,VOL=SER=PROD01,DISP=OLD</pre> <p>If the VOL/VOLG operands are specified on PRINT statements, any volume will be dynamically allocated if they are online, so the DISKxxxx DDs are usually not required.</p>
TAPExxxx DD STATEMENT(S)	Specifies the input backup data set to be used for the PRINT TVTOC statement. xxxx is any combination of valid alphanumeric digits (0-9, A-Z). The user may specify multiple TAPExxxx DD statements, and all of them will be processed in one execution. If the backup data set is not cataloged, the DD statement must specify the full data set name, volume serial, unit, file sequence number and disposition.
SORTLIB DD STATEMENT	Specifies the data set that contains the SORT functional modules. Required by selected PRINT functions, if your installation needs a SORTLIB DD Statement to perform sort operations.
SYSOUT DD STATEMENT	Specifies a data set for messages from the SORT program. Required by selected PRINT functions, if your installation's SORT program writes messages to a SYSOUT data set.
SORTWKnn DD STATEMENTS	Specify SORT work area data sets. Required by selected PRINT functions if your installation's SORT program required external work areas.
SYSIN DD STATEMENT	Specifies the control statement data set. Normally required for all ABR function unless input data is entered through PARM= on the execute statement.

53.03 FDRABRP ARCHIVE REPORT

PRINT P	ARCHIVE ,ADATE=yyddd yyyyddd ,COPY=1 2 ,DELETE=YES NO ,DSN=(dsname,...,dsname) ,DSG=(dsgroup,...,dsgroup) ,EXPIRE=YES NO ,FORMAT=PRT CRT ,HEX ,LINECNT=nn ,RECALL=YES NO	,REORG%=nn ,RESTORED=YES NO ,SDATE=yyddd yyyyddd ,SDAYS=nnnn ,SELTERR=YES NO ,SORT=YES NO ,SUM=YES NO ,VOL=(vvvvvv,...,vvvvvv) ,VOLG=(vvvvv,...,vvvvv) ,XDAYS=nnn
--------------------	---	--

PRINT ARCHIVE STATEMENT The report generated by the PRINT ARCHIVE statement contains information stored by ABR for data sets that have been archived, including: Data set name, DASD volume ID and type, date archived, data set attributes, backup file number and volume serial number(s) and when the backup will expire.

NOTE: If you specify a data set name with the DSN operand and this data set was archived multiple times, ABR will only display the most current copy. If you want to see all of the times the data set was ARCHIVED, specify DSG= the full data set name.

ICF VSAM ICF VSAM files are reported by the cluster name followed by its individual components. Selection criteria for VSAM files is only by cluster name.

OPERANDS

ADATE=	Specifies the date the data set was archived as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are assumed to be in the 21st century, e.g., 02123=2002.123. For readability, a period may be inserted between the year and day (e.g., 1997.123). ABR will only process the data sets which match this date.
COPY=	Specifies which backup copy number (1 or 2) to process for tests such as SDAYS and EXPIRE=.
	The default is that both copies will be tested.
DELETE=	YES selects only the archived data sets that have been marked for deletion (by the DELETE statement of FDRARCH) and are subject to removal from the ARCHIVE control file. NO selects only the archived data sets that have not been marked for deletion. The default is that the delete flag is not tested.

53.03 CONTINUED

DSN=	<p>Specifies one or more data set names from 1 to 44 characters in length. Only archived data sets or clusters having a name that match one of the names will be selected.</p> <p>NOTE: FDRABRP does not support data set name masking (patterns). For similar reports using masking, see FDREPORT in Section 54.</p>
DSG=	<p>Specifies one or more data set name prefixes from 1 to 44 characters in length. Only archived data sets or clusters having a name that begins with one or more of the prefixes will be selected.</p> <p>There is a special form of the DSG operand. Leading periods (.) after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed.</p> <p>EXAMPLE: DSG=..TST will select any data set with a third index level starting with 'TST'.</p> <p>NOTE: DSN and DSG operands may be repeated and/or intermixed. A total of up to 200 strings may be specified.</p>
EXPIRE=	<p>YES selects ONLY the archived data sets that have expired (data sets that are subject to removal from the ARCHIVE control file because they are past their expiration date). Expiration dates recorded as 99.000, 99.365 or 99.366 will be interpreted as "never expire".</p> <p>NO selects only the archived data sets that have not expired.</p> <p>The default is no expiration date check is made.</p>
FORMAT=	<p>PRT produces a 121 character line formatted for a printer.</p> <p>CRT produces a 79 character line formatted for a TSO terminal.</p> <p>The default is CRT if executing under TSO and PRT otherwise.</p>
HEX	<p>Specifies that the ARCHIVE control file is to be printed in a dump format (character/hexadecimal). Individual fields are NOT identified.</p> <p>The default is that the report is produced in character format, with the data identified by field within the records.</p>
LINECNT=	<p>Specifies the maximum number of lines each report page can contain, from 10 to 99, inclusive.</p> <p>The default is each page will contain a maximum of 58 lines.</p>
RECALL=	<p>YES selects only the archived data sets that were archived by FDRABR with the RECALL=YES option for auto-recall will be selected.</p> <p>NO selects only the archived data sets that have not been marked for auto-recall.</p> <p>The default is that the recall flag is not tested.</p>
REORG%=	<p>Specifies a minimum percent of free space the user requires in the ARCHIVE control file. If the free space is less than the nn%, message FDR428 is issued, listing the actual percent of free space.</p> <p>The default is a minimum 10% free space in the ARCHIVE control file.</p>
RESTORED=	<p>YES selects only the archived data sets that have been restored by FDRABR will be selected.</p> <p>NO selects only the archived data sets that have not been restored.</p> <p>The default is that the restored flag is not tested.</p>

CONTINUED . . .

53.03 CONTINUED

- RESTORED=** **YES** selects only the archived data sets that have been restored by FDRABR will be selected.
NO selects only the archived data sets that have not been restored.
 The default is that the restored flag is not tested.
- SDATE=** Specifies the date the data set was archived as a Julian date (year plus day number) which may be in the form "yyyyddd" (e.g., 1997123) or "yyddd" (e.g., 97123). If the 2-digit year is used (yyddd), year numbers less than 70 are assumed to be in the 21st century, e.g., 02123=2002.123. For readability, a period may be inserted between the year and day (e.g., 1997.123). ABR will only process the data sets which were archived on or after this date. This option can reduce the execution time of the print program if many data sets have been archived.
 The default is that the entire ARCHIVE control file will be searched.
- SDAYS=** Specifies a value in days used to calculate a prior Julian date (today's date minus the SDAYS= value). This date is used as described above by SDATE operand. This option can reduce the execution time of the print program if many data sets have been archived.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no archived data sets that match the selection criteria on this PRINT command.
NO – Specifies that a condition code of 12 will not be set if there are no archived data sets that match the selection criteria on this PRINT command.
 The default is YES.
- SORT=** **YES** results in the ARCHIVE report being printed in data set name sequence. [See Section 53.02](#) for possible JCL modifications required for sorting.
NO results in the report being printed in archive date sequence (most current to least current).
 The default is SORT=YES except when the input device or the output device is a TSO terminal, which causes SORT=NO to become the default.
- SUM=** **YES** results in the printing of the summary report after the detail listing.
NO results in the summary report being bypassed.
 The default is YES.
- VOL=** Specifies the DASD volume serial number which the data sets were archived. Multiple volume serial numbers may be specified in parentheses separated by commas.
- VOLG=** Specifies the prefix of the DASD volume which the data sets were archived. Multiple volume groups may be specified in parentheses separated by commas.
NOTE: VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified.
- XDAYS=** Specifies a value in days used to calculate a future expiration date (today's date plus the XDAYS= value). This date is used when calculating the number of expired entries printed in the summary report. This value is also used if the user selects expired entries by coding EXPIRE=YES operand.
 The default is 10 days.

CONTINUED . . .

53.03 CONTINUED

A B R A R C H I V E R E P O R T

FDR303 CARD IMAGE - * PRINT ARCHIVE,SDAYS=90

GENERATE ABR ARCHIVE REPORT *

FDR400 AUTOMATIC BACKUP/RECOVERY REPORT - FDRABRP VER 5.3/01P -INNOVATION DATA PROCESSING DATE: 1995.363 PAGE - 12

FDRABR ARCHIVE REPORT --- DSNNAME SEQUENCE

ARCHIVE CONTROL FILE IS -- IDPLB.Archive.PROD

AND IS LOCATED ON DASD VOLUME -- IDPLB3

***--DATA SET NAME----	*-DSK VOL-*	*--DATE OF--*	D/S REC	BLOCK	LBP	BK	***--DSN	TAPE	INFO----
SERIAL S TY ARCHIVE EXPDT ORG FM SIZE LRECL	ALLO	R/TRK*	TY CN SUFFIX	FN VOL-SERIAL(S)					
USTEST.T04500.COPY1. G0005V00	IDPLB4 1 0E 95.341	96.006 PS VB	32760 32756	1T	0	0E 1 B195341A	0 IDPBK0		
USTEST.T04500.COPY1. G0006V00	IDPLB4 1 0E 95.341	96.006 PS FB	12004 12000	300T	0	0E 1 B195341A	0 IDPBK0		
USTPROD.BKLEVEL.PAT. D950911.T092748	IDPBK0 1 0E 95.278	99.365 PS VB	32760 32756	2729T	2728	80 2 B295341A	7 BA0077		
USTPROD.BKLEVEL.PC245D. D951002.T175555	IDPBK0 1 0E 95.278	99.365 PS VB	32760 32756	1239T	1238	80 1 B195278A	101 BA0074,BA0075		
USTSYS.ST220A.ASM	IDPLB2 1 0E 95.334	95.364 PO FB	19040 80	9C	95	0E 1 B195334A	0 IDPBK0		
USTSYS.V230.ASM	IDPLB0 1 0E 95.334	95.364 PO FB	9040 80	89T	88	0E 1 B195334A	0 IDPBK0		

***--VOL TY FIELD SHOWS DEVICE TYPE OF DATA SET ARCHIVED--*

08--2314	09--3330	0D--3330-11	0A--3340	0B--3350	06--2305-01	07--2305-02	0E--3380	0C--3375
0E--3380-E	0E--3380-K	85--F6421	0E--3380-1	0E--3380-2	0F--3390	0F--3390-2	0F--3390-F	0E--3380-E4
0E--3380-3	0F--3390-3	0F--3390-9	04--9345-1	04--9345-2	0E--3380-K4	0F--3390-E2	0E--3380-K5	09--FAKE01
0F--3995-53								

***--BK TY FIELD SHOWS DEVICE TYPE OF ARCHIVE BACKUP--*

0E--3380 80--3480

***** INFORMATION FOR FIELD LABELED *DSN SUFFIX*

THE TAPE DSNNAME IS CONSTRUCTED AS FOLLOWS: THE FIRST INDEX IS FDABR. THE SECOND IS THE DISK SERIAL NUMBER PRECEDED BY A V THEN THE THIRD IS FOUND UNDER DSN-SUFFIX EXAMPLE FOR DATA SET NAME USTSYS.V230.ASM (FDRABR.VIDPLB0.B295334A)

***CLUSTER - DEFINES A VSAM ICF BASE CLUSTER FOLLOWED BY ITS ASSOCIATED COMPONENTS

***** INFORMATION FOR FIELD LABELED R/TRK*

WHEN THE DATA IS FOLLOWED BY A ASTERISK(*). THIS INDICATES THE FIELD VALUE IS THE RELATIVE TRACK ADDRESS OF THE BEGINNING OF THE DATA SET. WHEN THE ASTERISK IS NOT PRESENT THE VALUE IS THE NUMBER OF USED TRACKS (LAST BLOCK POINTER).

*****--ALLOC TYPE-----*

C - CYLINDER ALLOCATION
T - TRACK ALLOCATION
A - ABSOLUTE TRACK ALLOCATION

THIS REPORT PROVIDES THE A B R USER A LIST OF ALL DATA SETS THAT HAVE BEEN ARCHIVED FROM DASD STORAGE. THE REPORT ALSO FURNISHES THE USER A LIST OF TAPE VOLUME(S) THAT CONTAIN THE DATA SETS THAT HAVE BEEN ARCHIVED FROM DASD STORAGE DEVICES.

S A M P L E A R C H I V E R E P O R T U N D E R T S O

*** LOCATED ON DASD VOLUME -- IDPLB3

DSN-ABRT.TESTZAP.CNTL	DVOL-IDPLB1 3380	VSEQ---1
ARCHIVE DATE-1995.334 BK/TY(1)-3380 BK/TY(2)-3480 EXT--1 NO/CYL-----14		
DSORG-PO RECFM-FB	BLKSIZE---6160 LRECL---80 RECALL	LBP-----162
TDSN(1)-FDRABR.VIDPLB1.B195334A EXPDT-1995.364	FIL---0	VOL-IDPBK0
TDSN(2)-FDRABR.VIDPLB1.B295334A EXPDT-1999.365	FIL-154	VOL-BA0076

DSN-ABRTEST.IAMTVSAM.SAVE	DVOL-IDPLB1 3380	VSEQ---1
ARCHIVE DATE-1995.334 BK/TY(1)-3380 BK/TY(2)-3480 EXT--1 NO/TRK-----3		
DSORG-PO RECFM-F	BLKSIZE---4096 LRECL--4096 RECALL	LBP-----1
TDSN(1)-FDRABR.VIDPLB1.B195334A EXPDT-1995.364	FIL---0	VOL-IDPBK0
TDSN(2)-FDRABR.VIDPLB1.B295334A EXPDT-1999.365	FIL-154	VOL-BA0076

CONTINUED . . .

PRINT	CATLG	,MAXGEN=nnnn
P	,BKDAYS=nn	,SELTERR=<u>YES</u> NO
	,FORMAT=PRT CRT	,VOL=(vvvvvv,...,vvvvvv)
	,LINECNT=nn	,VOLG=(vvvvv,...,vvvvv)
	,MAXCYC=nn	

OPERANDS	<p>BKDAY= If the most recent backup for a DASD volume processed by this command was taken more than "nn" days ago, the FDR429 warning message is printed, to warn you that this volume may be accidentally bypassed during ABR backup processing. You can override the return code 12 from the FDR429 message by specifying FDR429=0 on the PRINT CATLG statement.</p> <p>The default is 7 days.</p>
	<p>FORMAT= PRT produces a 121 character line formatted for a printer. CRT produces a 79 character line formatted for a TSO terminal.</p> <p>The default is CRT if executing under TSO and PRT otherwise.</p>
	<p>LINECNT= Specifies the maximum number of lines each report page report can contain, from 10 to 99, inclusive.</p> <p>The default is each page will contain a maximum of 58 lines.</p>
	<p>MAXCYC= Specifies the number of cycles you wish to print within a single generation of ABR backup tapes. The order of selection will be the highest cycle number to the lowest number.</p> <p>The default is that ALL cycles will be printed.</p>
	<p>MAXGEN= Specifies the number of generations you want to print, starting from the most current.</p> <p>The default is that ALL generations will be printed.</p>
	<p>SELTERR= YES – Specifies that a condition code of 12 will be set if there are no entries in the ABR backup catalog that match the selection criteria on this PRINT command. NO – Specifies that a condition code of 12 will not be set if there are no entries in the ABR backup catalog that match the selection criteria on this PRINT command.</p> <p>The default is YES.</p>
	<p>VOL= Specifies the DASD volume serial number whose backups are to be printed. Multiple volume serial numbers may be specified in parentheses separated by commas.</p>

53.04 CONTINUED

VOLG=

Specifies the prefix of the DASD volume whose backups are to be printed.
Multiple volume groups may be specified in parentheses separated by commas.

If neither VOL= nor VOLG= is specified, all volumes whose backups are recorded in the ABR catalog will be printed.

NOTE: VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified.

VOLUME BACKUP REPORT

FDR303 CARD IMAGE - * PRINT CATLG

GENERATE VOLUME BACKUP REPORT *

FDR400 AUTOMATIC BACKUP/RECOVERY REPORT - FDRABRP VER 5.2/40T -INNOVATION DATA PROCESSING DATE: 1995.363 PAGE- 3

FDRABR VOLUME BACKUP REPORT

VOLSER	GEN	CYCLE	TYPE	DUMP DATE	TAPE FILE DATA SET NAME	COPY	FILE	TAPE VOLUME(S)	REQUIRED TO RESTORE
IDPLB0	396	03	DSF	1995.340	FDRABR.VIDPLB0.C1039603	1	24	B80332	
		04	DSF	1995.341	FDRABR.VIDPLB0.C1039604	1	30	B80337	
	395	00	FDR	1995.326	FDRABR.VIDPLB0.C1039500	1	1	B80223	
		01	DSF	1995.331	FDRABR.VIDPLB0.C1039501	1	1	B80269	
		02	DSF	1995.331	FDRABR.VIDPLB0.C1039502	1	2	B80269	
		03	DSF	1995.332	FDRABR.VIDPLB0.C1039503	1	13	B80273	
		04	DSF	1995.333	FDRABR.VIDPLB0.C1039504	1	29	B80278	
		05	DSF	1995.334	FDRABR.VIDPLB0.C1039505	1	34	B80283	
IDPLB1	429	00	FDR	1995.362	FDRABR.VIDPLB1.C1042900	1	4	B80486,B80489	
		00	FDR	1995.362	FDRABR.VIDPLB1.C2042900	2	4	BV1004,BV1007	
SYSLB2	86	00	FDR	1995.362	FDRABR.VSYSLB2.C1008600	1	5	B80489,B80491,B80494,B80496	
		00	FDR	1995.362	FDRABR.VSYSLB2.C2008600	2	5	BV1007,BV1010,BV1012,BV1016	

FDR499 FDRABRP PROCESSING COMPLETED

THE ABOVE REPORT PROVIDES THE USER WITH A B R BACKUP INFORMATION.
THIS LIST SEPARATES ALL DISK PACKS INTO GROUPS WHICH INCLUDE
ALL GENERATIONS AND CYCLES KEPT IN THE ABR CATALOG FOR ALL
DISK PACK(S) INITIALIZED FOR ABR PROCESSING.
THIS GIVES THE USER A HARDCOPY OF BACKUP DATA SET NAMES AND THEIR
TAPE VOLUME SERIAL NUMBERS.

SAMPLE CATALOG REPORT FROM TSO

VOLSER-IDPLB0	GENERATION--395	CYCLE---04	BKPD-1995.333	TYPE DUMP---DSF
	TAPE DSN-FDRABR.VIDPLB0.C1039504		COPY---1	FILE NO---29
	TAPE VOL-B80278			
VOLSER-IDPLB0	GENERATION--395	CYCLE---05	BKPD-1995.334	TYPE DUMP---DSF
	TAPE DSN-FDRABR.VIDPLB0.C1039505		COPY---1	FILE NO---34
	TAPE VOL-B80283			
VOLSER-IDPLB1	GENERATION--429	CYCLE---00	BKPD-1995.362	TYPE DUMP---FDR
	TAPE DSN-FDRABR.VIDPLB1.C1042900		COPY---1	FILE NO---4
	TAPE VOL-B80486,B80489			
VOLSER-IDPLB1	GENERATION--429	CYCLE---00	BKPD-1995.362	TYPE DUMP---FDR
	TAPE DSN-FDRABR.VIDPLB1.C2042900		COPY---2	FILE NO---4
	TAPE VOL-BV1004,BV1007			

CONTINUED . . .

53.05 FDRABRP SCRATCH REPORT

```

PRINT          SCRATCH
P
               ,DSN=(dsname,...,dsname)
               ,DSG=(dsgroup,...,dsgroup)

               ,FORMAT=PRT|CRT

               ,LINECNT=nn

               ,OLDBACKUP=CUR|ALL|(n..n)

               ,SELTERR=YES|NO

               ,VOL=(vvvvvv,...,vvvvvv)
               ,VOLG=(vvvvv,...,vvvvv)

               ,XREF

```

The report generated by the PRINT SCRATCH statement gives you critical information that has been captured by the ABR DADSM Pre-processing exit for data sets that have been scratched or renamed, including: Data set name, DASD volume ID and type, generation and cycle number, backup date, backup file number and volume serial numbers.

ICF VSAM ICF VSAM files are reported by cluster name. The individual components are not recorded in the SCRATCH catalog.

OPERANDS

DSN= Specifies one or more data set names from 1 to 44 characters in length. Only scratched data sets or clusters having a name that match one of the names will be selected.

NOTE: FDRABRP does not support data set name masking (patterns). For similar reports using masking, see FDREPORT in [Section 54](#).

DSG= Specifies one or more data set name prefixes from 1 to 44 characters in length. Only scratched data sets or clusters having a name that begins with one or more of the prefixes will be selected.

There is a special form of the DSG operand. Leading periods (.) after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed.

EXAMPLE: DSG=..TST will select any data set with a third index level starting with 'TST'.

NOTE: DSN and DSG operands may be repeated and/or intermixed. A total of up to 200 strings may be specified.

FORMAT= **PRT** produces a 121 character line formatted for a printer.
CRT produces a 79 character line formatted for a TSO terminal.

The default is CRT if executing under TSO and PRT otherwise.

LINECNT= Specifies the maximum number of lines each report page can contain, from 10 to 99, inclusive.

The default is each page will contain a maximum of 58 lines.

53.05 CONTINUED

- OLDBACKUP=** Specifies the relative backup numbers that are to be printed when processing entries from the SCRATCH catalog.
ALL requests all backups associated with the data set to be printed.
CUR requests the current backup only to be printed.
A numeric value of 0 through 3 permits the user to select relative backups for printing, where 0 is the current backup and 3 is the oldest.
The default is CUR.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no entries in the ABR SCRATCH catalog that match the selection criteria on this PRINT command.
NO – Specifies that a condition code of 12 will not be set if there are no entries in the ABR SCRATCH catalog that match the selection criteria on this PRINT command.
The default is YES.
- VOL=** Specifies the DASD volume serial number from which the data sets were scratched. Multiple volume serial numbers may be specified in parentheses separated by commas.
- VOLG=** Specifies the prefix of the DASD volume from which the data sets were scratched. Multiple volume groups may be specified in parentheses separated by commas.
NOTE: VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified.
- XREF** Specifies that the report is to include backup information, including the backup data set name, file number, and tape volume serial numbers.
The default is that only the ABR generation and cycle numbers of the backup is printed.

CONTINUED . . .

53.05 CONTINUED

A B R S C R A T C H R E P O R T

FDR303 CARD IMAGE - * PRINT SCRATCH,XREF,OLDBACKUP=ALL

GENERATE ABR SCRATCH REPORT *

FDR400 AUTOMATIC BACKUP/RECOVERY REPORT - FDRABRP VER 5.3/01P -INNOVATION DATA PROCESSING DATE: 1995.363 PAGE- 1

FDRABR LIST OF SCRATCHED DATA SETS

---DATA SET NAME---	*--DISK SERIAL	VOLUME--* TYPE	SEQ	GEN	CYC	BK NO	BACKUP DATE	TAPE DSN SUFFIX	TAPE FILE	TAPE VOLUME(S)	REQUIRED TO RESTORE
ABRT.TESTZAP.CNTL	IDPLB1	3380	001	425	0	00	1995.326	C1042500	0003	B80225,B80231,B80234	
				424	0	01	*****	C1042400	****	CAUTION -- OLD BKUP NOT CATALOGED	
ABRTEST.IAMTVSAM.SAVE	IDPLB1	3380	001	425	0	00	1995.326	C1042500	0003	B80225,B80231,B80234	
				424	0	01	*****	C1042400	****	CAUTION -- OLD BKUP NOT CATALOGED	
AGK.SPFLOG1.LIST	TSOWK0	3380	001	207	1	00	1995.352	C1020701	0007	B80452	
AGK.SPFTEMP0.CNTL	TSOWK0	3380	001	207	0	00	1995.349	C1020700	0005	B80443	
				206	2	01	1995.346	C1020602	0017	B80403	
AGK.UPSTREAM.BACKUPS. G0002V00	IDPWK0	3380	001	164	1	00	1995.352	C1016401	0009	B80452	
AGK.UPSTREAM.BACKUPS. G0039V00	IDPWK0	3380	001	160	4	00	1995.320	C1016004	0035	B80210	
AGK.UPSTREAM.BACKUPS. G0040V00	IDPWK0	3380	001	160	4	00	1995.320	C1016004	0035	B80210	
AGK.UPSTREAM.BACKUPS. G0041V00	IDPWK0	3380	001	160	4	00	1995.320	C1016004	0035	B80210	
USTSYS.V230.ASM	IDPLB0	3380	001	395	0	00	1995.326	C1039500	0001	B80223	
				394	0	01	*****	C1039400	****	CAUTION -- OLD BKUP NOT CATALOGED	

FDR499 FDRABRP PROCESSING COMPLETED

S A M P L E S C R A T C H R E P O R T F R O M T S O

DSN-AGK.SPFLOG1.LIST	VOL-TSOWK0	3380	VSEQ---1
*****-----BACKUP INFORMATION-----*****			
BKD(00)-1995.352 SFX-C1020701 FN-0007 VOLS-B80452			
DSN-AGK.SPFTEMP0.CNTL	VOL-TSOWK0	3380	VSEQ---1
*****-----BACKUP INFORMATION-----*****			
BKD(00)-1995.349 SFX-C1020700 FN-0005 VOLS-B80443			
BKD(01)-1995.346 SFX-C1020602 FN-0017 VOLS-B80403			
DSN-AGK.UPSTREAM.BACKUPS.G0002V00	VOL-IDPWK0	3380	VSEQ---1
*****-----BACKUP INFORMATION-----*****			
BKD(00)-1995.352 SFX-C1016401 FN-0009 VOLS-B80452			
DSN-AGK.UPSTREAM.BACKUPS.G0039V00	VOL-IDPWK0	3380	VSEQ---1
*****-----BACKUP INFORMATION-----*****			
BKD(00)-1995.320 SFX-C1016004 FN-0035 VOLS-B80210			

CONTINUED . . .

53.06 FDRABRP VOLUME STATUS REPORT

```

PRINT          VOLSTAT
P
               ,BKDAYS=nnn

               ,FORMAT=PRT|CRT

               ,INFO=YES|NO

               ,LINECNT=nn

               ,SELTERR=YES|NO

               ,VOL=(vvvvvv,...,vvvvvv)
               ,VOLG=(vvvvvv,...,vvvvvv)

```

The report generated by the PRINT VOLSTAT statement provides current information related to ABR processing for DASD volumes including: DASD volume serial, last backup date, current generation and cycle number, tape retention period, type of backup last processed, expiration date of the current generation, and informative messages.

OPERANDS	BKDAYS=	<p>If the most recent backup for a DASD volume processed by this command was taken more than "nn" days ago, the FDR429 warning message is printed, to warn you that this volume may be accidentally bypassed during ABR backup processing.</p> <p>The default is 7 days.</p>
	FORMAT=	<p>PRT produces a 121 character line formatted for a printer. CRT produces a 79 character line formatted for a TSO terminal.</p> <p>The default is CRT if executing under TSO and PRT otherwise.</p>
	INFO=	<p>NO – specifies that ABR is not to print the information messages for the volumes being selected. Example of an information message is an indication of whether the volume is enabled for ARCHIVE or SUPERSCRATCH.</p> <p>Default is YES, all information messages are printed.</p>
	LINECNT=	<p>Specifies the maximum number of lines each report page can contain, from 10 to 99, inclusive.</p> <p>The default is each page will contain a maximum of 58 lines.</p>
	SELTERR=	<p>YES – Specifies that a condition code of 12 will be set if there are no online volumes that match the selection criteria on this PRINT command. NO – Specifies that a condition code of 12 will not be set if there are no online volumes that match the selection criteria on this PRINT command.</p> <p>The default is YES.</p>
	VOL=	<p>Specifies the DASD volume serial number to be reported. Multiple volume serial numbers may be specified in parentheses separated by commas.</p>

53.06 CONTINUED

VOLG=

Specifies the prefix of the DASD volume to be reported. Multiple volume groups may be specified in parentheses separated by commas.

If neither VOL= nor VOLG= is specified, all online DASD volumes will be reported.

NOTE: VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified.

CONTINUED . . .

53.06 CONTINUED

VOLUME STATUS REPORT

FDR303 CARD IMAGE - * P VOLSTAT

```

FDR400  FDRABR VOLUME STATUS UTILITY - FDRABRVC VER 5.3/01P - INNOVATION DATA PROCESSING  DATE- 1996.162  PAGE- 0001
VOLUME  LAST  BACKUP  MAXIMUM  CURRENT AUT HITR RETP1 BKEXPDT1
SERIAL  DATE    TYPE   GEN  CYC  GEN  CYC  CYC  LOTR RETP2 BKEXPDT2 COMMENTS
-----
IDPWK2 1996.159  FDR    5  50   15   0   0  80%   35 1996.194 * INFO ONLY -- ARCHIVE FUNCTION DISABLED FOR THE VOLUME *
                                           50%      0      SCRATCH FUNCTION DISABLED FOR THE VOLUME *
                                           OLD BACKUP TABLE ENABLED FOR THE VOLUME **
                                           THRESHOLD VALUES - HIGH 80% -- LOW 50% *
SYSLB3 1996.159  FDR    5  50   76   0   0  80%   35 1996.194 * INFO ONLY -- ARCHIVE FUNCTION DISABLED FOR THE VOLUME *
                                           50%      0      SCRATCH FUNCTION DISABLED FOR THE VOLUME *
                                           OLD BACKUP TABLE ENABLED FOR THE VOLUME **
                                           THRESHOLD VALUES - HIGH 80% -- LOW 50% *

```

SAMPLE VOLUME STATUS REPORT FROM TSO

```

VOLSER---IDPWK2
  LAST ABR DATE--1996.159  RETPD(1)----35  MAX GEN-----5  MAX CYCLE-----50
  LAST BACKUP TYPE---FDR  RETPD(2)-----0  CUR GEN-----15  CUR CYCLE----0/00
  BACKUP(1) EXPIRES---1996.194
  STATUS---* INFO ONLY -- ARCHIVE FUNCTION DISABLED FOR THE VOLUME *
  STATUS---* INFO ONLY -- SCRATCH FUNCTION DISABLED FOR THE VOLUME *
  STATUS---* INFO ONLY -- OLD BACKUP TABLE ENABLED FOR THE VOLUME **
  STATUS---* INFO ONLY -- THRESHOLD VALUES - HIGH 80% -- LOW 50% *

VOLSER---SYSLB3
  LAST ABR DATE--1996.159  RETPD(1)----35  MAX GEN-----5  MAX CYCLE-----50
  LAST BACKUP TYPE---FDR  RETPD(2)-----0  CUR GEN-----76  CUR CYCLE----0/00
  BACKUP(1) EXPIRES---1996.194
  STATUS---* INFO ONLY -- ARCHIVE FUNCTION DISABLED FOR THE VOLUME *
  STATUS---* INFO ONLY -- SCRATCH FUNCTION DISABLED FOR THE VOLUME *
  STATUS---* INFO ONLY -- OLD BACKUP TABLE ENABLED FOR THE VOLUME **
  STATUS---* INFO ONLY -- THRESHOLD VALUES - HIGH 80% -- LOW 50% *

```

CONTINUED . . .

53.07 FDLABRP VTOC REPORT

PRINT P	VTOC	,FORMAT=PRT CRT
	,AGE=nnn	,LINECNT=nn
	,AGEINC=nnn	,ONLINE
	,COMBINE	,SELTERR= <u>YES</u> NO
	,DETAIL= <u>YES</u> NO	,SUMPFX=YES ALL
	,DSN=(dsname,...,dsname)	,VOL=(vvvvvv,...,vvvvvv)
	,DSG=(dsgroup,...,dsgroup)	,VOLG=(vvvvv,...,vvvvv)

The report generated by the Print VTOC statement provides current information about existing data sets taken from the VTOCs of specific volumes, including: data set name, DASD volume serial, last ABR cycle, volume sequence number, last reference date, special ABR indicators, data set attributes and allocation information. The SUMPFX option allows summarization by high-level index.

OPERANDS	AGE=	<p>If the summary report is printed, AGE= will specify the number of days (1 to 999) since a data set has been referenced as used by the aging summary for the VTOC Report.</p> <p>The default is 30.</p>
	AGEINC=	<p>If the summary report is printed, AGEINC= will specify the number of days (1 to 999) added to the value to derive the next control break in the aging summary for the VTOC Report. The summary report will display four aging breaks.</p> <p>The default is 30.</p>
	COMBINE	<p>Specifies that the report will be produced in data set/volume serial number sequence, merging the various disk volumes into one combined report.</p> <p>NOTE: The use of COMBINE requires that the execution Job Control include SORT-related DD statements as described in Section 53.02.</p> <p>The default is that each volume processed will have its own separate report.</p>
	DETAIL=	<p>NO – results in ONLY a track utilization and data set aging summary to be printed. The SUM=YES option must also be active otherwise no reports will be output. The reporting of DSCB information related to each data set is bypassed.</p> <p>The default is DETAIL=YES.</p>
	DSN=	<p>Specifies one or more data set names from 1 to 44 characters in length. Only data sets or clusters having a name that match one of the names will be selected</p> <p>NOTE: FDLABRP does not support data set name masking (patterns). For similar reports using masking, see FDREPORT in Section 54.</p>

53.07 CONTINUED

- DSG=** Specifies one or more data set name prefixes from 1 to 44 characters in length. Only data sets or clusters having a name that begins with one or more of the prefixes will be selected.
- There is a special form of the DSG operand. Leading periods (.) after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed.
- EXAMPLE: DSG=..TST will select any data set with a third index level starting with 'TST'.
- NOTE:** DSN and DSG operands may be repeated and/or intermixed. A total of up to 200 strings may be specified.
- FORMAT=** **PRT** produces a 121 character line formatted for a printer.
CRT produces a 79 character line formatted for a TSO terminal.
- The default is CRT if executing under TSO and PRT otherwise.
- LINECNT=** Specifies the maximum number of lines each report page can contain, from 28 to 99 inclusive.
- The default is each page will contain a maximum of 58 lines.
- ONLINE** Causes all online DASD volumes to be included in the report.
- If neither ONLINE nor VOL/VOLG is specified, the default is that only disk volumes pointed to by DISKxxxx DD statements will be included.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no data sets that match the selection criteria on this PRINT command.
NO – Specifies that a condition code of 12 will not be set if there are no data sets that match the selection criteria on this PRINT command.
- The default is YES.
- SUM=** NO suppresses the track utilization and data set aging summary. Only the information from the DSCBs for the data sets on the volumes selected will be printed. You must insure that the DETAIL=YES option is active.
- The default is SUM=YES.
- SUMPFX=** Specifies that the report on ABRSUM summarize by high-level data set index.
YES will produce only a report by index level.
ALL also produce AGING information for the higher data sets by index level.
- VOL=** Specifies string(s) from 1 to 6 characters in length. All volumes online to the system that match the selection criteria (i.e.: the compare length is six) will be selected for printing.
- The default is deferred to VOLG operand.
- VOLG=** Specifies string(s) from 1 to 6 characters in length. All volumes online to the system that match the selection criteria (i.e.: the compare length is the length of data specified) will be selected for printing.
- If neither ONLINE nor VOL/VOLG is specified, the default is that only disk volumes pointed to by DISKxxxx DD statements will be included.
- NOTE:** VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified. The volumes specified will be dynamically allocated.

CONTINUED . . .

53.07 CONTINUED

SAMPLE VTOC REPORT

FDR303 CARD IMAGE - * P VTOC,VOL=IDPLB2

FDR400 AUTOMATIC BACKUP/RECOVERY VTOC LIST - FDRABRV VER 5.3/01P - INNOVATION DATA PROCESSING DATE- 1996.162 PAGE- 1

VTOC LIST OF VOLUME SERIAL NUMBER IDPLB5 - DEVICE TYPE 3380-K

DATA SET NAME	SERIAL	LS	CY	SEQ	DATE	IND	ORG	RECFM	SIZE	LRECL	ALLOC	FREE	CT	CCC-HH	CCC-HH
****VTOC	IDPLB5			001	0000.000				0	0	15	0	1	01327-00	01327-14
***ABR MODEL--LAST GEN=0417	IDPLB5	00	001	1996.159					0	0	0	0	0		
FDRABR.VIDPLB5															
ABRARC.GDGDSN01	IDPLB5	00	001	1990.274		**	FB		2000	100	0	0	0		
ABRARC.GDG01	IDPLB5	00	001	1990.274		**	FB		2000	100	0	0	0		
***CLUSTER---BAB.BAB.DDIR															
BAB.BAB.DDIR.DATA	IDPLB5	00	001	1996.120		EF	U		4096	0	135	0	1	02307-00	02315-14
BAB.BAB.DDIR.INDEX	IDPLB5	00	001	1996.120		EF	U		4096	0	1	0	1	00426-11	00426-11
BGR.P510400.BKUP	IDPLB5	00	001	1996.130		PS	U		32760	0	29	0	1	02288-00	02289-13
BGR.TEST.CLIST	IDPLB5	00	001	1996.162		PO	FB		3120	80	6	1	2	00426-12	00427-00
														00422-04	00422-05

ABR INDICATORS --

A=ALWAYS BACKUP/NEVER ARCHIVE R=ARCHIVING REQUESTED C=NO CURRENT BACKUP
 N=NORMAL BACKUP/NEVER ARCHIVE U=UPDATE INDICATED X=EXCLUDE FROM ABR
 Z=ABR INDICATORS CONTAMINATED

IN SUMMARY - VOLUME IDPLB5 --- DEVICE TYPE 3380-K
 *** TRACK UTILIZATION *** PER *** DATA SET STATISTICS ***

TOTAL TRACKS-----39825
 ALLOCATED TRACKS-----38680 98% TOTAL DATA SETS-----562
 PS/PO ALLOCATED---24222 PS/PO DATA SETS-----384
 PS/PO USED-----18810
 PS/PO UNUSED---5412 22%
 OTHER ALLOCATED---14457 OTHER DATA SETS-----178
 FREE TRACKS-----1145 2% TOTAL EXTENTS-----945

*** DATA SET AGING SUMMARY -- BASED ON LAST REFERENCED DATE **

LREFDATE	DAYS SINCE	DATA	ALLOC	USED	UNUSED
REFERENCED	SETS	TRACKS	TRACKS	TRACKS	TRACKS
1996.132	1 - 30	219	20793	16800	3993
1996.102	31 - 60	41	2197	1824	373
1996.072	61 - 90	55	5258	4977	281
****.***	91 - ***	247	10431	9666	765

FDR426 PROCESSING OF VOLUME IDPLB5 COMPLETED

SAMPLE VTOC REPORT

DSN-ABRARC.GDGDSN01 VOL-IDPLB5 3380-K VSEQ---1
 LAST REF--1990.274 LAST CYC---0 EXT COUNT---0 ALLOC-----0 FREE-----0
 DSORG-----** RECFM---FB BLKSIZE--2000 LRECL---100

SAMPLE VTOC REPORT FROM TSO

DSN-ABRARC.GDG01 VOL-IDPLB5 3380-K VSEQ---1
 LAST REF--1990.274 LAST CYC---0 EXT COUNT---0 ALLOC-----0 FREE-----0
 DSORG-----** RECFM---FB BLKSIZE--2000 LRECL---100

CLUSTER---BAB.BAB.DDIR
 DSN-BAB.BAB.DDIR.DATA VOL-IDPLB5 3380-K VSEQ---1
 LAST REF--1996.120 LAST CYC---0 EXT COUNT---1 ALLOC-----135 FREE-----0
 DSORG-----EF RECFM---U BLKSIZE--4096 LRECL-----0
 EXTENT-----0 BEGINNING CCHH-----02307-00 ENDING CCHH-----02315-14

DSN-BAB.BAB.DDIR.INDEX VOL-IDPLB5 3380-K VSEQ---1
 LAST REF--1996.120 LAST CYC---0 EXT COUNT---1 ALLOC-----1 FREE-----0
 DSORG-----EF RECFM---U BLKSIZE--4096 LRECL-----0
 EXTENT-----0 BEGINNING CCHH-----00426-11 ENDING CCHH-----00426-11

CONTINUED . . .

53.08 FDRABRP DATA SET BACKUP REPORT

```

PRINT
P

      BACKUP

      ,COMBINE

      ,DSN=(dsname,...,dsname)
      ,DSG=(dsgroup,...,dsgroup)

      ,FORMAT=PRT|CRT

      ,LINECNT=nn

      ,OLDBACKUP=CUR|ALL|(n...n)

      ,ONLINE

      ,SELTERR=YES|NO

      ,VOL=(vvvvvv,...,vvvvvv)
      ,VOLG=(vvvvvv,...,vvvvvv)

```

The report generated by the PRINT BACKUP statement is similar to the PRINT VTOC report, except that includes current ABR backup information for each data set. If OLDBACKUP is ENABLED for the volume, ABR will optionally print previous backup information.

OPERANDS	COMBINE	<p>Specifies that the report will be produced in data set/volume serial number sequence, merging the various disk volumes into one combined report.</p> <p>NOTE: The use of COMBINE requires that the execution Job Control include SORT-related DD statements as described in Section 53.02.</p> <p>The default is that each volume processed will have its own separate report.</p>
	DSN=	<p>Specifies one or more data set names from 1 to 44 characters in length. Only data sets or clusters having a name that match one of the names will be selected</p> <p>NOTE: FDRABRP does not support data set name masking (patterns). For similar reports using masking, see FDREPORT in Section 54.</p>
	DSG=	<p>Specifies one or more data set name prefixes from 1 to 44 characters in length. Only data sets or clusters having a name that begins with one or more of the prefixes will be selected.</p> <p>There is a special form of the DSG operand. Leading periods (.) after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed.</p> <p>EXAMPLE: DSG=..TST will select any data set with a third index level starting with 'TST'.</p> <p>NOTE: DSN and DSG operands may be repeated and/or intermixed. A total of up to 200 strings may be specified.</p>
	FORMAT=	<p>PRT produces a 121 character line formatted for a printer.</p> <p>CRT produces a 79 character line formatted for a TSO terminal.</p> <p>The default is CRT if executing under TSO and PRT otherwise.</p>

53.08 CONTINUED

- LINECNT=** Specifies the maximum number of lines each report page can contain, from 28 to 99 inclusive.
The default is each page will contain a maximum of 58 lines.
- OLDBACKUP=** Specifies the relative oldbackup number(s) that is(are) to be printed.
ALL requests all backups associated with the data set is to be printed.
CUR requests that only the current backup information is to be printed.
A numeric value of 0 through 13 permits the user to select relative backup tapes for printing (0 is the most recent backup and 13 the oldest).
The default is 'CUR'.
- ONLINE** Causes all online DASD volumes to be included in the report.
If neither ONLINE nor VOL/VOLG is specified, the default is that only disk volumes pointed to by DISKxxxx DD statements will be included.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no data sets that match the selection criteria on this PRINT command.
NO – Specifies that a condition code of 12 will not be set if there are no data sets that match the selection criteria on this PRINT command.
The default is YES.
- VOL=** Specifies string(s) from 1 to 6 characters in length. All volumes online to the system that match the selection criteria (i.e.: the compare length is six) will be selected for printing.
The default is deferred to VOLG operand.
- VOLG=** Specifies string(s) from 1 to 6 characters in length. All volumes online to the system that match the selection criteria (i.e.: the compare length is the length of data specified) will be selected for printing.
If neither ONLINE nor VOL/VOLG is specified, the default is that only disk volumes pointed to by DISKxxxx DD statements will be included.
NOTE: VOL and VOLG operands may be repeated and/or intermixed. A total of up to 100 strings may be specified. The volumes specified will be dynamically allocated.

CONTINUED . . .

53.08 CONTINUED

SAMPLE BACKUP REPORT

FDR303 CARD IMAGE - * P BACKUP,VOL=IDPLB5,DSG=BAB.,OLDBACKUP=ALL

FDR400 FDR400 AUTOMATIC BACKUP/RECOVERY REPORT - FDRABRP VER 5.3/01 - INNOVATION DATA PROCESSING DATE: 1998.008

FDR303 CARD IMAGE - * PRINT BACKUP,DSG=BAB,VOL=IDPLB5,OLDBACKUP=ALL *

FDR491 BACKUP REPORT FUNCTION STARTED - 16.22.48

FDR400 AUTOMATIC BACKUP/RECOVERY VTOC LIST - FDRABRV VER 5.3/01 - INNOVATION DATA PROCESSING DATE- 1998.008 PAGE- 1

DSNAME/XREF LIST FROM VOLUME IDPLB5 - DEVICE TYPE 3380-K

VOLUME VL D/S TRACKS LAST REF BACKUP BK TAPE DSN TAPE

DATA SET NAME SERIAL SQ ORG ALLOC FREE DATE DATE NO SUFFIX FILE TAPE VOLUMES REQ TO RESTORE

BAB.ISPF4.PROFILE IDPLB5 01 PO 12 1 1997.364 1998.002 00 C1045800 0003 B90064

1997.364 01 C1045716 0143 B90062

1997.346 02 C1045705 0134 B90060

1997.339 03 C1045700 0005 B90058

1997.338 04 C1045617 0150 B90028

1997.337 05 C1045616 0126 B90028

1997.336 06 C1045615 0102 B90028

1997.335 07 C1045614 0078 B90028

1997.330 08 C1045613 0054 B90028

1997.329 09 C1045612 0030 B90028

1997.328 10 C1045611 0006 B90028

1997.325 11 C1045610 0237 B90057

1997.324 12 C1045609 0213 B90057

1997.323 13 C1045608 0189 B90054

BAB.TEST.LOAD IDPLB5 01 PO 30 11 1997.339 1998.002 00 C1045800 0003 B90064

1997.346 01 C1045705 0134 B90060

1997.339 02 C1045700 0005 B90058

1997.335 03 C1045614 0078 B90028

1997.328 04 C1045611 0006 B90028

1997.323 05 C1045608 0189 B90054

1997.322 06 C1045607 0165 B90054

1997.311 07 C1045600 0005 B90054

1997.307 08 C1045519 0159 B90052

1997.301 09 C1045516 0095 B90052

1997.276 10 C1045500 0005 B90048

1997.261 11 C1045409 0194 B90042

1997.248 12 C1045400 0005 B90042

1997.234 13 C1045315 0065 B90040

ASTERISK(*) AFTER LAST REFERENCED DATE DENOTES UPDATE INDICATOR

FDR426 PROCESSING OF VOLUME IDPLB5 COMPLETED

FDR499 FDRABRV (5.3/01) PROCESSING COMPLETED

FDR492 BACKUP REPORT FUNCTION ENDED - 16.23.11

FDR400 AUTOMATIC BACKUP/RECOVERY REPORT - FDRABRP VER 5.3/01 -INNOVATION DATA PROCESSING DATE: 1998.008 PAGE- 2

FDR999 FDRABRP (5.3/01) PROCESSING COMPLETED

SAMPLE BACKUP REPORT UNDER TSO

DSN-BAB.TEST.LOAD VOL-IDPLB5 3380-K VSEQ---1
 LAST REF--1996.162 DSORG----PO ALLOC-----60 FREE-----7
 ***** BACKUP INFORMATION *****
 BKD(00)-1996.159 SFX-C1041700 FN-002 VOLS-B80532,B80536,B80540
 BKD(01)-1996.156 SFX-C1041602 FN-017 VOLS-B80514
 BKD(02)-1996.155 SFX-C1041601 FN-004 VOLS-B80507
 BKD(03)-1996.152 SFX-C1041600 FN-003 VOLS-B80468,B80470,B80474
 BKD(04)-1996.151 SFX-C1041503 FN-019 VOLS-B80457,B80465
 BKD(05)-1996.150 SFX-C1041502 FN-017 VOLS-B80451,B80458
 BKD(06)-1996.145 SFX-C1041500 FN-002 VOLS-B80397,B80400,B80405
 BKD(07)-1996.143 SFX-C1041403 FN-023 VOLS-B80364
 BKD(08)-1996.141 SFX-C1041401 FN-003 VOLS-B80351
 BKD(09)-1996.138 SFX-C1041400 FN-002 VOLS-B80303,B80311,B80320
 BKD(10)-1996.137 SFX-C1041304 FN-035 VOLS-B80293
 BKD(11)-1996.136 SFX-C1041303 FN-021 VOLS-B80283
 BKD(12)-1996.135 SFX-C1041302 FN-016 VOLS-B80279
 BKD(13)-1996.134 SFX-C1041301 FN-004 VOLS-B80271

CONTINUED . . .

53.09 FDRABRP TAPE VTOC REPORT

The report generated by the PRINT TVTOC command gives the user access to VTOC information and/or VVDS information from the control records at the beginning of a dump tape created by FDR, FDRDSF, or FDRABR. The type of report is selected by using the LIST= and/or VSAM= operands.

The input tape is identified by a DD statement with a DDNAME prefixed by 'TAPE'. If multiple TAPExxxx DD statements are present, all of them will be processed, unless FRDD= is specified. A VSAMPRT SYSOUT DD statement is required if you wish to generate a listing of the VVDS information.

```

PRINT          TVTOC
P

,DSN=(dsname,...,dsname)
,DSG=(dsgroup,...,dsgroup)

,FRDD=ddname

,LINECNT=nn

,LIST=ABR|VTOC|DUMP|NO

,SELTERR=YES|NO

,VSAM=YES|NO|DUMP

```

The report generated by the PRINT TVTOC statement provides VTOC information and/or VVDS information from the control records at the beginning of a backup created by FDR, FDRDSF, or FDRABR, so that you can see the details of the data sets included in that backup. The type of report is selected by using the LIST= and/or VSAM= operands.

The input backup data set is identified by a DD statement with a DDNAME prefixed by 'TAPE'. If multiple TAPExxxx DD statements are present, all of them will be processed, unless FRDD= is specified. A VSAMPRT SYSOUT DD statement is required if you wish to print the VVDS information.

OPERANDS	DSN=	Specifies one or more data set names from 1 to 44 characters in length. Only data sets or clusters in the backup having a name that match one of the names will be selected
	DSG=	<p>Specifies one or more data set name prefixes 1 to 44 characters in length. Only data sets or clusters in the backup having a name that begins with one or more of the prefixes will be selected.</p> <p>There is a special form of the DSG operand. Leading periods (.) after DSG= indicate that the group name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed.</p> <p>EXAMPLE: DSG=..TST will select any data set with a third index level starting with 'TST'.</p> <p>NOTE: DSN and DSG operands may be repeated and/or intermixed. A total of up to 200 strings may be specified.</p>
	FRDD=	<p>Specifies the name of the DD statement from which to read the backup. This allows a selective execution when the JCL contains multiple TAPExxxx DD statements.</p> <p>The default is that all TAPExxxx DD statements will be processed.</p>

53.09 CONTINUED

- LINECNT=** Specifies the maximum number of lines each report page can contain, from 28 to 99, inclusive.
- The default is that each page will contain a maximum of 58 lines.
- LIST=** Specifies the format of the listing of VTOC information.
- ABR** – the listing of VTOC information will be in a format similar to that of program FDRABRP, command PRINT VTOC, FORMAT=PRT.
- VTOC** – the listing of VTOC information will be in a format similar to that of program IEHLIST, command LISTVTOC FORMAT.
- DUMP** – the listing of VTOC information will be in a format similar to that of program IEHLIST, command LISTVTOC DUMP.
- NO** – NO listing of VTOC information will be produced.
- Default is LIST=ABR.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no data sets backed up to the tapes being processed that match the selection criteria on this PRINT command.
- NO** – Specifies that a condition code of 12 will not be set if there are no data sets backed up to the tapes being processed that match the selection criteria on this PRINT command.
- The default is YES.
- VSAM=** Specifies the format of the listing of VVDS information for ICF VSAM data sets. The VSAMPRT DD statement must be present if this option is specified as other than NO.
- YES** – the listing of VVDS information will be in a format similar to that of program IDCAMS, command LISTCAT ALL. ABR will fill in as much of this report as is possible from the BACKUP. Some fields, such as Paths for alternate indexes and password protection, are not available from the backup tape.
- DUMP** – the listing of VVDS information will be dump format for each VVR, with hexadecimal on the left and EBCDIC on the right.
- NO** – NO listing of VVDS information will be produced.
- The default is VSAM=YES if VSAMPRT is present, VSAM=NO if not.

CONTINUED . . .

53.09 CONTINUED

SAMPLE TVTOC REPORT IN IEHLIST FORMAT

FDR303 CARD IMAGE - * P TVTOC,LIST=VTOC

```

FDR400 AUTOMATIC BACKUP/RECOVERY VTOC - FDRABRP VER 5.3/01P - INNOVATION DATA PROCESSING DATE: 1996.162 PAGE - 1
FORMAT 4 DSCB NO AVAIL/MAX DSCB /MAX DIRECT NO AVAIL NEXT ALT FORMAT 6 LAST FMT 1 VTOC EXTENT THIS DSCB
VI DSCBS PER TRK BLK PER TRK ALT TRK TRK(C-H) (C-H-R) DSCB(C-H-R)/LOW(C-H) HIGH(C-H) (C-H-R)
89 13302 50 45 15 3339 0 0 0 0 20 0 50 0 1 20 0 0 1 1
-----UNABLE TO CALCULATE EMPTY SPACE.
-----DATA SET NAME----- ID SER NO SEQ NO CREDIT EXPDT NO EXT DSORG RECFM OPTCD BLKSIZE
***ABR MODEL--LAST GEN=0067 FDRABR.VS93004 1 S93004 1 04493 36099 0 00 00 0
LRECL KEYLEN INITIAL ALLOC 2ND ALLOC/LAST BLK PTR(T-R-L) USED PDS BYTES FMT 2 OR 3(C-H-R)/DSCB(C-H-R)
0 0 TRKS 0 0 0 58786 50 0 0 0 0 1 5
***--THIS DATASET HAS NO EXTENTS--**
-----DATA SET NAME----- ID SER NO SEQ NO CREDIT EXPDT NO EXT DSORG RECFM OPTCD BLKSIZE
ADPBQ0.ISPF.ISPPROF 1 S93004 1 29790 00000 1 PO FB 00 3120
LRECL KEYLEN INITIAL ALLOC 2ND ALLOC/LAST BLK PTR(T-R-L) USED PDS BYTES FMT 2 OR 3(C-H-R)/DSCB(C-H-R)
80 0 TRKS 1 0 3 55522 0 0 0 0 3 3 13
EXTENTS NO LOW(C-H) HIGH(C-H)
0 2361 0 2361 0
-----ON THE ABOVE DATA SET THERE ARE 0 EMPTY TRACK(S).

```

SAMPLE TVTOC VSAM REPORT

```

DATA ----- APLDSNU.DSNDBD.Z0DGEN.ECDMVSRU.I0001.A001
IN-CAT --- DSNCAT2
HISTORY
DATASET-OWNER----- (N/A) CREATION----- (N/A)
RELEASE----- (N/A) EXPIRATION----- (N/A)
PROTECTION-PSWD----- (N/A) RACF----- (N/A)
ASSOCIATIONS
CLUSTER--APLDSNU.DSNDBC.Z0DGEN.ECDMVSRU.I0001.A001
ATTRIBUTES
KEYLEN-----0 AVGLRECL-----0 BUFSIZE-----8192 CISIZE-----4096
RKP-----0 MAXLRECL-----0 EXCPXIT----- (NULL) CI/CA-----72
SHROPTNS(3,3) UNIQUE ERASE NOWRITECHK NOREPLICAT REUSE NONSPANNED RECOVERY
UNORDERED NOIMBED
STATISTICS
REC-TOTAL-----0 SPLITS-CI-----0 EXCPS-----0
REC-DELETED-----0 SPLITS-CA-----0 EXTENTS-----3
REC-INSERTED-----0 FREESPACE-%CI-----0 SYSTEM-TIMESTAMP:
REC-UPDATED-----0 FREESPACE-%CA-----0 X'0000000000000000'
REC-RETRIEVED-----0 FREESPC-BYTES-----0
ALLOCATION
SPACE-TYPE-----TRACK HI-ALLOC-RBA-----3538944
SPACE-PRI-----60 HI-USED-RBA-----3244032
SPACE-SEC-----6
VOLUME
VOLSER-----S93004 PHYREC-SIZE-----4096 HI-ALLOC-RBA-----3538944 EXTENT-NUMBER-----3
DEVTYPE-----3390-3 PHYRECS/TRK-----12 HI-USED-RBA-----3244032 EXTENT-TYPE-----X'40'
VOLFLAG-----PRIME TRACKS/CA-----6
EXTENTS:
LOW-CCHH----X'033D0001' LOW-RBA-----0 TRACKS-----60
HIGH-CCHH----X'03410000' HIGH-RBA-----2949119
LOW-CCHH----X'02FD000E' LOW-RBA-----2949120 TRACKS-----6
HIGH-CCHH----X'02FD0004' HIGH-RBA-----3244031
LOW-CCHH----X'02FD0005' LOW-RBA-----3244032 TRACKS-----6
HIGH-CCHH----X'02FD000A' HIGH-RBA-----3538943

```

CONTINUED . . .

53.10 FDRABRP JCL EXAMPLES

The following examples illustrate the most common ways of executing ABR REPORT functions. Note that, for convenience, all STEPLIB/JOBLIB DD statements have been omitted in the examples; they may be required, depending on your installation's placement of ABR.

ISPF DIALOG SUPPORT Most of these reports are available under the ABR ISPF dialogs, on the FDRABR REPORT PANEL (ISPF option A.1).

ABR CATALOG REPORT Produce a report detailing all backup activity that has been recorded in the ABR catalog. The report consists of information pertaining to the disk volume dumped, type and date of backup and information pertaining to the tape used for the dump function.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//SYSIN      DD    *
PRINT       CATLG
```

ABR SCRATCH REPORT Produce a report detailing all data sets that have been scratched or renamed and recorded in the ABR scratch catalog. The report consists of information pertaining to the DASD volume and type of device that the data set existed on, and the backup tape file and volumes where the data was stored. The current backup information and old backups if any, will be displayed.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//SYSIN      DD    *
PRINT       SCRATCH,XREF,OLDBACKUP=ALL
```

SUMMARY BY INDEX Produce a report detailing the utilization of DASD space by USERID (highest level index). The number of data sets owned by each ID, the tracks occupied and the types of data sets by organization are detailed. Aging statistics by last reference date are also printed.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//ABRSUM     DD    SYSOUT=*
//SYSIN      DD    *
PRINT       VTOC,SUMPFX=ALL,ONLINE,DETAIL=NO
```

TAPE VTOC REPORT This example illustrates the JCL required to produce a report detailing information from the formatted VTOC at beginning of FDR, DSF and ABR backup tapes. The report is displayed in ascending data set name sequence for each 'TAPExxxx' DD statement.

Note: This report is available to all FDR users.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//VSAMPRT    DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//TAPE0001   DD    DISP=SHR,DSN=FDRABR.VMVSR51.C1001000
//           DD    VOL=SER=100233,UNIT=TAPE,LABEL=4
//SYSIN      DD    *
PRINT       TVTOC,LIST=ABR                      ABR FORMAT VTOC REPORT
PRINT       TVTOC,LIST=DUMP                      IEHLIST DUMP OF VTOC
```

53.10 CONTINUED**ARCHIVE
REPORT**

Produce a report detailing all data sets that have been archived by ABR. The report consists of information pertaining to the disk volume from which the data sets were archived and the date they were archived. Also, certain data set attributes and information pertaining to the tape used for the dump function are displayed. The ARCHIVE DD statement points to the ARCHIVE Control File. Since SORT=YES is the default in batch, to display information in data set name order, SORT DD statements may be required.

```
//PRINT      EXEC   PGM=FDRABRP
//SYSPRINT   DD     SYSOUT=*
//SYSOUT     DD     SYSOUT=*
//ABRMAP     DD     SYSOUT=*
//ARCHIVE    DD     DSN=FDRABR.ARCHIVE,DISP=SHR
//SORTLIB    DD     DSN=SYS1.SORTLIB,DISP=SHR
//SORTWK01   DD     SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SORTWK02   DD     SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SORTWK03   DD     SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SYSIN      DD     *
PRINT       ARCHIVE
```

**ARCHIVE
REPORT BY
DATE**

Produce a report detailing all data sets that have been archived by ABR on a given date. The ARCHIVE Control File whose name is in the FDR Global Option Table will be dynamically allocated. SORT=NO is specified so that the data sets are not sorted by name; they will be displayed in the order they were archived.

```
//PRINT      EXEC   PGM=FDRABRP
//SYSPRINT   DD     SYSOUT=*
//SYSOUT     DD     SYSOUT=*
//ABRMAP     DD     SYSOUT=*
//SYSIN      DD     *
PRINT       ARCHIVE,SORT=NO,ADATE=1997.140
```

**ARCHIVE
EXPIRATION
PREDICTION**

Produce a report detailing all archived data sets which will expire in the next 15 days. EXPIRE=YES limits the report to expired data sets but XDAYS=15 includes data sets which will expire within 15 days as well.

```
//PRINT      EXEC   PGM=FDRABRP
//SYSPRINT   DD     SYSOUT=*
//SYSOUT     DD     SYSOUT=*
//ABRMAP     DD     SYSOUT=*
//SYSIN      DD     *
PRINT       ARCHIVE,SORT=NO,EXPIRE=YES,XDAYS=15
```

VTOC REPORT

Produce a report detailing information of the volumes referenced by the VOL parameter. The report is displayed in ascending data set name sequence for each volume specified.

```
//PRINT      EXEC   PGM=FDRABRP
//SYSPRINT   DD     SYSOUT=*
//ABRMAP     DD     SYSOUT=*
//SYSIN      DD     *
PRINT       VTOC,VOL=(PACK01,PACK02)
```

CONTINUED . . .

53.10 CONTINUED

BACKUP REPORT Produce a report detailing information for all data sets in the VTOCs of the volumes referenced by a DISKxxxx DD statement, in addition to information related to the most current backup tape on which the data was stored. The report is displayed in ascending data set name sequence for each volume specified.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//DISK1      DD    UNIT=DASD ,VOL=SER=PACK01 ,DISP=SHR
//DISK2      DD    UNIT=DASD ,VOL=SER=PACK02 ,DISP=SHR
//SYSIN      DD    *
PRINT       BACKUP
```

COMBINED VTOC REPORT Produce a report detailing information on all data sets in the VTOCs of all volumes referenced by the VOL parameter. The report is displayed in ascending data set name sequence for all volumes specified. Since COMBINE implies sorting, SORT DD statements may be required.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//SYSOUT     DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//SORTLIB    DD    DSN=SYS1.SORTLIB ,DISP=SHR
//SORTWK01   DD    SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SORTWK02   DD    SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SORTWK03   DD    SPACE=(TRK,50,,CONTIG),UNIT=DASD
//SYSIN      DD    *
PRINT       VTOC ,COMBINE ,VOL=(PACK01,PACK02)
```

ABR VOLUME REPORT Produce a report detailing the ABR status of all disk volumes online to the system.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//SYSIN      DD    *
PRINT       VOLSTAT
```

MULTIPLE REPORTS Produce an Archive Report, a Catalog Report, and a Volume Status Report in a single job step. The Archive Report will show the archived data sets in the reverse order of archival (the most recently archived data sets will appear first), and will only include information for data sets that were archived from volumes PACK01 and PACK02. The Catalog Report will only include information for backups of volumes PACK01 and PACK02.

```
//PRINT      EXEC  PGM=FDRABRP
//SYSPRINT   DD    SYSOUT=*
//ABRMAP     DD    SYSOUT=*
//ARCHIVE    DD    DSN=FDRABR.ARCHIVE ,DISP=SHR
//SYSIN      DD    *
PRINT       ARCHIVE ,SORT=NO ,VOL=(PACK01,PACK02)
PRINT       CATLG ,VOL=(PACK01,PACK02)
PRINT       VOLSTAT ,VOL=(PACK01,PACK02)
```

CONTINUED . . .

THIS PAGE WAS INTENTIONALLY LEFT BLANK

54.01 GENERALIZED REPORT WRITER (FDREPORT)

FDREPORT is a generalized report writer which can generate reports on the DASD-related data that you specify, in a format that you specify. You can generate simple reports with very simple control statements, or complex reports using the great power of FDREPORT.

You can generate reports on a few data sets, large numbers of data sets, whole DASD volumes, or your entire installation. You can report on live data, or data archived or backed up by ABR. You can select data through the system catalogs, or directly from disk volumes or other sources. You can generate a data file for further reporting or pass the data to other programs for further analysis. You can even punch JCL or control statements for other programs using the report data.

FDREPORT is designed to report on large amounts of data from various sources without sacrificing performance. FDREPORT will require less elapsed time and system resources (CPU time and EXCPs) for a given function than any other competitive product.

NOTE: for YEAR2000 support, FDREPORT is capable of displaying all date fields with 4-digit year numbers. This is specified by the DATEFORMAT= operand. Prior to V5.3, the default was DATEFORMAT=YYDDD (2-digit year); in V5.3 and above, the default is DATEFORMAT=YYYYDDD (4-digit year). This may affect the format of existing FDREPORT jobstreams if they depend on the length of the default date format. See [Section 54.10](#) (the DEFAULT statement) for instructions on permanently changing the default. Of course, the date format can be overridden in any particular FDREPORT job step.

**INPUT
SOURCES**

FDREPORT gathers input from a variety of sources:

- MVS catalogs - a data set name filter can be specified to quickly search catalogs for the required data sets
- VTOC of DASD volumes - information from the DSCBs of the data sets is extracted
- VVDS of DASD volumes - for VSAM clusters and SMS-managed data sets
- ABR backup information - for data sets processed by ABR incremental backups
- Archive/Application Control File (ACF) - for archived data sets and those processed by application backups
- DFHSM or DFSMSHsm MCDS or BCDS data
- data recorded by previous FDREPORT executions
- backup data sets created by ABR, FDR, FDRDSF, or SAR, for information on DASD data sets in those backups

From the hundreds of different fields available, each relating to an attribute of a data set or a volume, you can select a subset to be reported, and can optionally specify the positioning of those fields in the report, although FDREPORT does most of the formatting work for you. In addition to simple printed reports, FDREPORT can sort and summarize the data selected. It can also generate control statements and/or JCL for other utilities based on the data collected. Although most reports are based on individual data sets, you can report on the state of entire DASD volumes (such as volume free space).

By default, FDREPORT does not attempt to serialize access to the input sources it is reading. Input data sets (such as the ACF, MCDS and BCDS) are allocated with DISP=SHR unless you provide a DD statement for them with DISP=OLD. VTOCs and VVDSs are read without any sort of ENQ or RESERVE unless you enable the RESERVE option ([see Section 54.30](#)) which allows updates to take place during the FDREPORT execution. In rare cases, a VTOC/VVDS update may cause FDREPORT to fail or generate incorrect output unless RESERVE is enabled.

54.01 CONTINUED . . .



At any computer installation, users at a variety of levels need to be able to refer to accurate and timely information about the use of DASD resources. Accurate information is the only basis for sound decision making and future planning. Without adequate information it is impossible to determine how efficiently DASD space is being used, or predict how needs will grow. Accurate information is also essential for solving existing problems and for preventing problems from arising in the future.

FDREPORT provides easy-to-read management level reporting that allows you to easily see how effectively your DASD storage is being used, broken down by departments, projects, or individual users.

Since storage costs in most shops represent a significant part of the data processing budget, the ability to accurately monitor and adjust for current and future needs will result in significant cost savings and a more competitive cost for your end user.

Innovation strongly recommends that you run the Innovation Health Check jobstreams documented in [Section 54.02](#). These will help you understand the power, flexibility and efficiency of FDREPORT.

**DATA SET
REPORTS**

Most reports are about individual data sets or ICF VSAM clusters. Although FDREPORT has a default report format (giving some basic information about each data set) and some "canned" report formats (similar to those produced by FDRABRP in [Section 53](#)), you will probably want to customize your report.

FDREPORT collects its data into hundreds of "fields", each containing some aspect of the data set, such as record format, size, CISIZE, creation date, and many, many others. [Section 54.31](#) contains a complete list of the fields available. You can select the data sets to be reported based on the values of most of these fields, and you can report on any set of them you select. You also control the positioning of the fields in the report.

**VOLUME
REPORTS**

FDREPORT can also report on the status of entire DASD volumes, selecting and reporting on fields such as device type, available space, number of data sets (by type) and many others. [Section 54.31](#) also contains the list of fields available for volume reports. You have the same selection and reporting options for volume reports as you do for data set reports.

SORTING

The report generated by FDREPORT can be sorted using most of the report fields. For example, you may request that it report on all of the data sets currently online in size order. FDREPORT will invoke the system SORT product when required, and can dynamically allocate required sort libraries and work areas.

SUMMARIES

FDREPORT can generate summaries on many of the report fields. These summaries can show you the various values that the summary fields had and counts of the occurrences of each value. The control break facility can cause summaries to be shown at various points in the report, when the value of some field changes.

PUNCHING

FDREPORT can generate control statements and JCL (or any arbitrary text) using a user-provided mask for the format of the data to be "punched", substituting the values of FDREPORT report fields into that mask. For example, FDREPORT can generate ABR control statements.

**DATA
EXTRACT**

FDREPORT can write the selected data to an extract file (in a unique FDREPORT format). The extract file can be used as input for further reports. This allows you to gather the data once and then report on it in various formats or using varying selection criteria.

It is possible to use this extract file as input to other data analysis programs, such as SAS, if they can read the extract file format. FDREPORT can also "print" data in a simple tabular format (no headings or page breaks) for input into other programs.

54.01 CONTINUED . . .

FDREPORT STATEMENTS The FDREPORT statements specify the format of the report, the sort sequence, summary requirements and which data sets or volumes are to be selected. **Note that a PRINT statement is always the last statement in any group of statements, since it causes the report to actually be generated. The statements are:**

DEFAULT	Processing defaults.
TITLE	User defined title line.
HEADING	User-defined replacement column heading lines.
SELECT	Selection criteria for the data sets to be included in the report.
EXCLUDE	Criteria for excluding certain data sets from the report.
XSELECT	Selection criteria, using more flexible specifications than the SELECT statement.
XEXCLUDE	Criteria for exclusion, using more flexible specifications than the EXCLUDE statement.
REPORT	Selects data fields to be printed.
SORT	Requests sorting on selected data fields.
SUMMARY	Selects data fields to be summarized with optional control break criteria.
BREAK	Criteria for control breaks and summaries.
PUNCH	Controls generation of control statements and/or JCL from report data
PRINT	Generates the report requested by the preceding statements.
CANCEL	Cancels preceding statement specifications, when producing multiple reports in one execution.
EXECUTE	Causes FDREPORT statements to be read from a library.

The statements are documented in the order shown. PRINT is the statement which causes a report to actually be generated. The statements that precede PRINT plus operands on PRINT itself define the report to be generated. Placing statements in the wrong order, such as SELECTs after PRINT, will usually result in an incorrect report. You may have multiple PRINT statements in a given FDREPORT input to generate multiple reports; note that statements will remain in effect for subsequent PRINTs unless overridden or cancelled by a CANCEL statement.

NOTE: FDREPORT contains features which are not documented in this manual, due to space limitations. These include special-purpose customization options, special modes of execution, field names and detailed information or special instructions about various selection and reporting criteria. For complete information about these, execute the following jobstream:

```
//HELP      EXEC  PGM=FDREPORT
//SYSPRINT   DD    SYSOUT=*
             HELP  ALL
```

This will print a lengthy document with additional information about FDREPORT. This report will fit on 8.5x11 inch paper if desired.

54.02 INNOVATION HEALTH CHECK

INNOVATION HEALTH CHECK Innovation has developed a series of jobstreams, based largely on FDREPORT, which can be used to report on the general status of your DASD installation, plus it also reports on various "problem" conditions, such as uncataloged data sets, and volumes running out of free space or room in the VTOC. The jobstreams are also good examples of both simple and sophisticated use of FDREPORT. These jobstreams are found in the FDR JCL library in members:

HCHECKI - index to the other members, similar to the list printed below.

HCHECK0 - FDREPORT Internal Parameter Check

HCHECK1 - DASD With Potential Problems:

- Disk Volumes More Than 80% Full
- Fragmentation Index, Worst First
- VTOCs More Than 80% Full
- VVDS More Than 80% Full
- VVDS In Multiple Extents
- VTOC Indexes More Than 80% Full
- Volume Mount & Use Status, SMS, VTOCIX Status
- List VVDSs and Check For Logical Errors
- VTOC Logical Errors

HCHECK2 - Reports For SMS Administrator:

- Volume SMS Status
- SMS Volumes With Disabled VTOC Indexes
- Dataset SMS Attributes
- Information On PDSE Data sets
- Uncataloged Datasets On SMS Volumes
- Non-Managed Datasets On SMS Volumes
- Datasets Ineligible For SMS Management

HCHECK3 - Reports For Performance Analyst:

- Disks With Potential Performance Problems
- Multi-Extent VSAM Data sets
- Multi-Extent Non-VSAM Data sets
- Datasets Likely To Get SX37 Abends
- Multi-Volume Data sets
- Poor VTOC/VVDS/VTOCIX Positions

HCHECK4 - Reports For Capacity Analyst:

- Disks With Potential Wasted Space
- Installed Dasd Summarised By Device Type
- Installed Dasd Summarised By Type and Model
- Overallocation In All Data sets
- Datasets Passed Their Expiration Dates
- Datasets Not Referenced in 60 Days
- Datasets Referenced In 60 Days, Overallocation
- Datasets With Inefficient Blocksizes by size
- Datasets With Inefficient Blocksizes by blocksize
- Small Datasets With Inefficient Blocksizes
- Total Space Available, Allocated, and Unused
- Volumes with More Than 45% Freespace
- VTOCs With More Than 45% Freespace
- VVDSs With More Than 45% Freespace
- VTOC Indexes With More Than 45% Freespace

54.02 CONTINUED . . .**HCHECK5** - Reports For VSAM Tuning Analyst:

- Space Occupied By Largest VSAM Files
- Busiest VSAM Files (EXCP Sort)
- Multi-Volume VSAM Files
- Clusters With High Inserts
- Multi-Extent Clusters
- Badly Split Clusters
- Clusters With 3 Or More Index Levels
- Overallocated VSAM Clusters

HCHECK6 - Reports Showing PDS Efficiency:

- Basic Stats On PDS Data sets
- Multi-Extent PDS Data sets
- PDS With Limited Freespace
- PDS With Limited Free Directory Blocks
- PDS Datasets With Excess Free Space
- Basic Stats On PDSE Data sets
- Space Reclaimable By PDS Compression

HCHECK7 - Report Data sets Which May Offend Standards:

- Uncataloged Data sets
- Wrongly Cataloged Data sets
- Datasets With Undefined DSORG
- Empty PS, PO, ICF-VSAM Data sets
- Unmovable Data sets
- Datasets Only In The Catalog

HCHECK8 - Project or Dasd Management Reports:

- Summary Of Data set Types
- Catalog Usage Report
- Summary Of ISAM Files
- Summary Of IAM Files
- Space Occupied By An Application
- Space Occupied By CNTL and JCL Files
- Space Occupied By LIST and OUTLIST Files
- Space Occupied By SYS1 & SYS2 Data sets
- Space Summarized By Index
- Space Summarized By 3rd Level
- DASD Utilization By Esoteric Unit Name SYSDA
- DASD Utilization By Esoteric Unit Name
- DASD Utilization By SMS Storage Group Name
- DASD Utilization By Volser Group
- DASD Utilization By Device Type

54.03 FDREPORT JCL REQUIREMENTS

The following Job Control Statements are required to execute the FDR Generalized Report Facility.

STEPLIB or JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDREPORT resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDREPORT) and region requirement (REGION=). Since the storage requirements of FDREPORT vary depending on the functions requested, Innovation recommends that you specify REGION=0M so that the maximum region is available.

SYSPRINT DD STATEMENT Specifies the primary output message data set; it is required. It is usually a SYSOUT data set. DCB characteristics are RECFM=FBA and LRECL=121; the blocksize will default to 1210 on disk or tape.

ABRMAP DD STATEMENT Specifies the report data set. It is usually a SYSOUT data set but can be directed to tape or disk. DCB characteristics are RECFM=FBA and LRECL of the value of the PAGEWIDTH operand (see [Section 54.19](#)) unless you override it; the default blocksize is BLKSIZE=0 (calculated by OPEN). When ABRMAP is not found within the JCL stream, the reports will be output to the SYSPRINT data set. The PAGEWIDTH operand will be ignored if ABRMAP is not present.

ABRSUM DD STATEMENT Specifies that the summary reports are to be printed on this data set. Usually a SYSOUT data set. If summaries are to be printed and ABRSUM is not found in the JCL, summaries will be printed on ABRMAP or on SYSPRINT if ABRMAP is also absent.

SYSUDUMP DD STATEMENT Recommended in all FDREPORT jobs in order to more easily diagnose error conditions which make FDREPORT abend. Usually a SYSOUT data set.

If you have the ABEND-AID product from COMPUWARE include the following so that a fully-formatted dump is produced: //ABNLIGNR DD DUMMY

SYSPUNCH DD STATEMENT Specifies the output data set when RPTYPE=SELPCH is specified, for the storage of generated statements. FDREPORT will force DCB characteristics of RECFM=FB,LRECL=80; any valid blocksize may be specified, but it will default to 80.

SYSPUNCH may be a DASD data set (sequential or a member of a PDS), but it may also be assigned to the JES internal reader "SYSOUT=(class,INTRDR)" to submit a complete jobstream for execution. You may also make SYSPUNCH a normal SYSOUT data set if you wish to view the generated statements.

DIAGPCH DD STATEMENT Specifies the output data set when ENABLE=(DIAGNOSEVVDS,PCHDIAGNOS) is specified, for the storage of a jobstream to correct VVDS errors. FDREPORT will force DCB characteristics of RECFM=FB,LRECL=80; any valid blocksize may be specified, but it will default to 80. DIAGPCH will usually be a DASD data set (sequential or a member of a PDS). After reviewing this jobstream, you can submit it to correct the VVDS errors.

SYSUT2 DD STATEMENT Required when the RPTYPE=DATA (or COMPDATA) or DATATYPE=EXTRACT operands are specified, for the storage of FDREPORT internal data records. Normally a DASD data set but it may be on tape. If you want to use this extract file as input to another reporting program (such as SAS), the format of the records is defined by a macro @RPTDS which is a member in the FDR ICL (Installation Control Library).

When RPTYPE=DATA or RPTYPE=COMPDATA is specified it will be an output data set and FDREPORT will force DCB characteristics of RECFM=VB,LRECL=8200; any valid blocksize can be specified, or FDREPORT will assign a default value.

If you execute several PRINT statements in the same FDREPORT step, each with RPTYPE=DATA (or COMPDATA), and you want the extracted data from each to accumulate in the SYSUT2 data set, you must either specify DISP=MOD in the JCL or the ENABLE=MODOUTPUT operand. Otherwise, only data from the last PRINT will be available.

For DATATYPE=EXTRACT, SYSUT2 is used as input; it must be a data file created by the

54.03 CONTINUED . . .

RPTYPE=DATA option of FDREPORT in an earlier step or job or earlier in this step. Although the format of the extract file has changed in various releases of FDR/ABR, FDREPORT can successfully process any extract file created by any prior release. You can concatenate several extract files as one input if they were created with RPTYPE=DATA; extract files created with RPTYPE=COMPDATA cannot be concatenated.

**SYSUT2 DD
STATEMENT
(continued)**

If required and not present in the JCL, FDREPORT will allocate an extract data set. If the EXTRACTDSN= operand is specified, the existing data set named will be allocated; you may optionally include the EXTRACTMEMBER= operand to select a member of a PDS for input or output. If EXTRACTDSN= is not included, a temporary data set will be allocated, which will be retained for the duration of the step so that it can be used as output and later as input when DATATYPE=EXTRACT is specified on another PRINT statement.

The SYSUT2 DD name may be changed by the DATADDNAME= operand.

**SYSUT1 DD
STATEMENT**

Required when the RPTYPE=DATA (or COMPDATA) and DATATYPE=EXTRACT operands are both specified on the same PRINT statement. Although SYSUT2 (described above) would normally be used for the data input, it cannot be used as both input and output in the same operation; in this case, SYSUT1 is used instead of SYSUT2 for the extracted data input, and SYSUT2 is used for output. The SYSUT1 DD name may be changed by the ALTDATADDNAME= operand.

**FDRLIB DD
STATEMENT**

Specifies a data set to be read for extra FDREPORT control statements. This data set is used if FDREPORT encounters an EXECUTE or PUNCH control statement. This data set can be a sequential data set or a partitioned data set but must have characteristics RECFM=FB and LRECL=80. The purpose of this data set is to store canned report formats and punch masks for execution by any user. The DD name may be changed by the FDRLIB= operand.

**DISKxxxx DD
STATEMENTS**

Optional and usually unnecessary. When the ENABLE=ONLINE operand is specified on PRINT, or VOL= or VOLG= operands are specified on statements preceding PRINT, the desired volumes will be dynamically allocated and DISKxxxx DD statements are not required.

If used, DISKxxxx DD statements specify the DASD volumes to be processed by statements that do not specify ONLINE, VOL= or VOLG=. The first four characters of the DDNAME must be 'DISK', and the remaining one to four characters may be any characters valid in a ddname (0-9, A-Z, \$, #, @). The DD statement should look like:

```
//DISKxxxx DD UNIT=unitname,VOL=SER=vvvvvvv,DISP=OLD
```

**TAPExxxx DD
STATEMENTS**

Required only if DATATYPE=TVTOC is specified. These DD statements must point to FDR-format backups on tape or disk; all backups created by FDR, FDRDSF, FDRABR and SAR can be used. All of the TAPExxxx DD statements in the step will be read when DATATYPE=TVTOC is used.

**ARCHIVE DD
STATEMENT**

Specifies the name of an ABR ARCHIVE control file, if the DATATYPE=ARCHIVE option is used to process data from an archive control file. This is not required if the ARCHIVE control file whose name is in the ARCDN option of the FDR Global Option Table is to be read; FDREPORT will dynamically allocate it under DD ARCHIVE#. You may also specify the ARCDN= operand of FDREPORT to dynamically allocate any ARCHIVE control file.

**MCDSDD and
BCDSDD DD
STATEMENTS**

(Optional) you can provide one or both of these DD statements if you use DATATYPE=MCDS or BCDS. They must point to the corresponding DFSMSHsm or DFHSM control file (MCDS or BCDS). If you provide the data set name of the HSM control file via the MCDSCLUSTER= and/or BCDSCLUSTER= operands, the DD statement is not necessary; FDREPORT will dynamically allocate the required file.

**SORT DD
STATEMENTS**

(Optional) If you include the SORT or generate a report with data from multiple disk volumes, you may need to specify DD statements (e.g. SORTLIB, SORTWKnn) required by your system SORT product.

However, any or all of these SORT DD statements may be omitted if you have requested dynamic allocation of the SORT data sets via the SORTALLOC= operand of PRINT.

54.03 CONTINUED . . .

SYSLIB DD STATEMENT Required if the CHANGE=PERM or CHANGE=RESET options are specified on a DEFAULT statement, in order to permanently change FDREPORT default processing options. Must point to the load library containing FDREPORT and you must have UPDATE authority to that library.

SYSIN DD STATEMENT Specifies the control statement data set. Usually a DD * or input data set.

FDREPORT under TSO FDREPORT may also be executed under TSO. Any required files must be preallocated (corresponding to the DD statements above). If the library containing ABR is in the linklist, then wherever you can enter a TSO command, you can simply type:

FDREPORT

if not, then you can type:

CALL 'fdrlibrary(FDREPORT)'

In either case, FDREPORT will prompt for statement input. Enter END to terminate FDREPORT. The SYSIN and SYSPRINT file names should be allocated to your terminal before invoking FDREPORT.

If you have the ABR ISPF panels installed, option A.1 (REPORTS) gives you the ability to generate pre-determined reports with FDREPORT (using the EXECUTE statement to select the report format from a library, with a SELECT statement generated from the panel.

However, ISPF option A.S (the SRS dialog, described in [Section 54.50](#)) internally uses FDREPORT to generate much of its data, giving you the ability to interactively do the same sort of selection and reporting as done by FDREPORT.

54.10 FDREPORT DEFAULT STATEMENT

DEFAULT	ALIASLEVEL=n	,LBPZERO= <u>INVALID</u> VALID
	,ALTDATADDNAME=ddname	,LINECNT=nn
	,ARCDSN=dsname	,MAXGDG=nnnn
	,BYTEFORMAT=BYTES <u>KILOBYTES</u> MEGABYTES	,MAXICF=nnnnnn
		,MAXONLINE=nnnn
	,CHANGE=PERM RESET <u>TEMP</u>	,MCDSCUSTER=mcscsclustname ,BCDSCUSTER=bcdscsclustname
	,COPY=BOTH <u>EITHER</u> 1 2	
	,DATADDNAME=ddname	,OLDBACKUP=(nn,nn,...) <u>ALL</u> <u>CUR</u>
	, DATATYPE=ARCHIVE BCDS CATALOG CATARCH CATVTOC EXTRACT MCDS TVTOC VOLDATA <u>VTOC</u>	PAGEWIDTH=nnnn
		,PCHDATEFORMAT=YYDDD MMDDY Y DDMMY MMDDYYYY DDMMYYYY Y YYYYDD
	,DATEFORMAT=YYDDD MMDDY DDMMY MMDDYYYY DDMMYYYY <u>YYYYDDD</u>	,RPTYPE=ABRVTOC ARCHIVE COMPDATA DATA <u>GENERATE</u> HEX NONE OSVTO C SELPCH TABLE XREF
	,DDCNT=nnn	
	,DISABLE=(option1,option2,...) ,ENABLE=(option1,option2,...)	,SKIP=n
	,DSKIP=n	,SORT=COMBINE <u>NO</u> YES
		,SORTALLOC=(option1,option2,...)
	,EXTRACTDSN=dsname	,SORTCORE=nnnnnnnn
	,EXTRACTMEMBER=membername	,SORTLIB=dsname
		,SORTMSGDDNAME=ddname
		,SORTPFX=cccc
	,FIELDPREFIX=c ,FIELDSUFFIX=c	,SUM=YES <u>NO</u> INDEX SUMDEVICE= <u>BASE</u> UNIQUE
	,FDRLIB=ddname	,SYSUTSPACE=nnnn ,SYSUTUNIT=unitname
	,FORMAT=CRT PRT	,TITLE=LEFT <u>CENTER</u> RIGHT
	,INDEXNUM=nn	,WORKDDNAMES=n ,WORKSPACE=nnnn ,WORKUNIT=unitname

DEFAULT STATEMENT The DEFAULT statement changes FDREPORT's default values for various options used when generating a report. Most of the operands on the DEFAULT statement are also operands of the PRINT statement, and can be specified there. The DEFAULT statement can be used when more than one report (more than one PRINT statement) is to be generated in one FDREPORT run; the

54.10 CONTINUED . . .

DEFAULT statement avoids having to specify options used in more than one report more than once. **Only the operands which are unique to the DEFAULT statement are described below; others are described in [section 54.19](#) under the PRINT statement.**

The DEFAULT statement may also be used to permanently update default values of most of the operands of DEFAULT in the ABR program library so that all subsequent executions of FDREPORT will use the new values without having to specify them. The underlined defaults shown above (and in the PRINT statement) are the distributed defaults; your installation may have changed them. The defaults will be permanently changed when the CHANGE=PERM or CHANGE=RESET operand is specified; this requires that a SYSLIB DD statement be included in the FDREPORT JCL pointing to the ABR program library. The current defaults may be displayed by the ENABLE=DISPLAY operand; this display also shows which values are eligible for permanent change by CHANGE=PERM, which can only be changed in the FDR Global Option Table, and which can not be changed.

OPERANDS	ARCDN=	<p>Specifies the data set name of the ARCHIVE Control File to be used as input if DATATYPE=ARCHIVE or DATATYPE=CATARCH is also specified and no ARCHIVE DD is present.</p> <p>The default is the ARCHIVE Control File name in the FDR Global Option Table and is usually FDRABR.ARCHIVE.</p>
	CHANGE=	<p>Allows the default values used by FDREPORT for other operands on the DEFAULT statement to be permanently changed. The available options are:</p> <p>PERM – permanently change the default values for many other operands specified on the DEFAULT statement, as described above.</p> <p>RESET – reset all defaults back to the original values distributed with FDREPORT.</p> <p>TEMP – operands on the DEFAULT statement affect this execution of FDREPORT only.</p> <p>PERM or RESET requires that a SYSLIB DD statement pointing to the ABR program library be included in the FDREPORT JCL; you must have UPDATE authority to this library.</p> <p>The default is TEMP.</p>
	DDCNT=	<p>Specifies the maximum number of user-specified DISKxxxx DD names that can be processed in any execution of FDREPORT, from 5 to 400. Note that Innovation recommends dynamic allocation of volumes to be processed, by use of the VOL= and ENABLE=ONLINE options.</p> <p>The default is 100.</p>
	FIELDPREFIX=	<p>Specifies the character that marks the beginning of a field name in a punch mask or TITLE statement. Valid characters are < &\$*%_>:#@ but the default of < should not be overridden unless you need to use that character as data in the mask.</p>
	FIELDSUFFIX=	<p>Specifies the character that marks the end of a field name in a punch mask or TITLE statement. Valid characters are < &\$*%_>:#@ plus a blank but the default of > should not be overridden unless you need to use that character as data in the mask.</p>

54.10 CONTINUED . . .

FDRLIB=	<p>Specifies the DD name to be used for reading FDRPT control statements by the EXECUTE statement or punch mask definitions by the PUNCH statement. Can also be specified on those statements.</p> <p>The default is FDRLIB.</p>
MAXICF=	<p>Specifies the maximum number of ICF VSAM clusters whose description can be retained in storage for any one report, from 100 to 10000. It is used only with DATATYPE=ARCHIVE or CATARCH.</p> <p>The default is 400 clusters.</p>
MCDSCLUSTER= BCDSCLUSTER=	<p>Specifies the cluster name of the DFHSM or DFSMSHsm MCDS or BCDS data set. The MCDS cluster will be dynamically allocated if DATATYPE=MCDS is used, and the BCDS cluster will be allocated if DATATYPE=BCDS.</p> <p>If these parameters are not specified, FDRPT expects that the JCL will contain a MCDSDD DD statement (for DATATYPE=MCDS) or a BCDSDD DD statement (for DATATYPE=BCDS).</p> <p>Only the operands which are unique to the DEFAULT statement are described above; others are described in section 54.19 under the PRINT statement.</p>

54.11 FDREPORT TITLE STATEMENT

TITLE**LINE='text'****,SKIP=n**

TITLE STATEMENT The TITLE statement defines a user-specified TITLE LINE to be displayed on every page of the report between the INNOVATION header identification and the data header lines. Under TSO, the INNOVATION header is not displayed. A maximum of one line may be specified.

OPERANDS **LINE=** Specifies the text to be printed or displayed. Must be enclosed in quotes. The number of characters specified must not exceed the page width. If the TITLE cannot be contained on one control statement (Column 1 to 71), the user can continue the text by specifying a '+' or '—' after the last character on this line.

If '+' is specified, FDREPORT will scan for the first non-blank character on the next input line, so you can start the continuation in any column. If '—' is specified FDREPORT will start with column 1 of the next input line.

EXAMPLE: TITLE LINE='BACKUP +
 REPORT'

The title text may contain FDREPORT fields, so that the title on each may contain data that related to the values displayed on that page. Any FDREPORT field name that is valid for SORT (see the tables in [Section 54.31](#)) may be included. You surround the field names with the FIELDPREFIX and FIELDSUFFIX characters that are in effect at the time that the TITLE statement is read (set by a previous DEFAULT or PUNCH statement; they default to < and >). FDREPORT will substitute the value for that field that is current at the time each new page is printed.

EXAMPLE: TITLE LINE='REPORT FOR VOLUME <VOL>'

SKIP= Specifies the number of lines to be left blank between the TITLE LINE and the data heading line, from 1 to 3.

Default is 1.

54.12 FDREPORT HEADING STATEMENT

HEADING **LINE(1)**='first heading line text'

 ,LINE(2)='second heading line text'

 ,LINE(3)='third heading line text'

HEADING STATEMENT By default, FDREPORT provides column headings with text which is descriptive of the field(s) in each column (e.g., DSNAME, SIZE). The HEADING statement allows you to specify replacement text of your choosing. It is your responsibility to line up the heading text with the actual columns generated by FDREPORT; this may take some experimentation.

OPERANDS **LINE(n)=** Specifies the replacement heading text for heading line n (n=1, 2, or 3).

 If the heading text cannot be contained in one control statement (columns 1 to 71), it may be continued using the same conventions described for the TITLE statement ([Section 54.11](#)).

54.13 FDREPORT (X)SELECT/(X)EXCLUDE STATEMENT

SELECT	DSN =(dsname,...,dsname)
EXCLUDE	DSG =(dsgroup,...,dsgroup)
	,VOL =(vvvvvv,...,vvvvvv)
	other FIELD names from the tables in Section 54.31
XSELECT	DSN op (dsname,...,dsname)
XEXCLUDE	DSG op (dsgroup,...,dsgroup)
	XDSN op (dsnamemask,...,dsnamemask)
	MEMNAMES op (member,...,member)
	,VOL op (vvvvvv,...,vvvvvv)
	,STORGRP op (ssssssss,...,ssssssss)
	,UNIT op (uuuuuu,...,uuuuuu)
	,UNITNAME op (uuuuuuuu,...,uuuuuuuu)
	other FIELD names from the tables in Section 54.31

SELECT/
XSELECT/
EXCLUDE/
XEXCLUDE
STATEMENTS

These statements act as a filter for the data sets to be processed by FDREPORT. The SELECT and EXCLUDE statements use fairly simple comparisons ("equal to" for some tests, "equal or greater than" for others), while the XSELECT and XEXCLUDE statements may use more sophisticated comparisons (less than or equal, greater than, not equal, etc.), as well as a data set name masking capability.

On the SELECT and EXCLUDE statements, operands may be followed ONLY by an equal sign (=). The value you provide after the equal will be compared to the value of the field associated with a given data set. The comparison will be true for certain fields if the value is exactly equal; for other fields it is true if the field is equal or greater than your value (see below).

On the XSELECT and XEXCLUDE statements, the operands may be followed by one of a number of comparison operators. Since one form of those operators involve special characters (such as the not(¬) and less-than(<)), alternate forms of each operator without special characters are provided. The operators are:

=	or	.EQ.	equal
¬ or ¬=	or	.NE.	not equal
<	or	.LT.	less than
>	or	.GT.	greater than
<=	or	.LE.	less than or equal to
>=	or	.GE.	greater than or equal to

The test will be true if the indicated comparison of the FIELD value and the value you provide is true. For example, SIZE>15,DSORG.NE.PS

54.13 CONTINUED . . .

On the XSELECT and XEXCLUDE statements, if the operation is equal (=, .EQ.) or not-equal (≠, .NE.), several values may be provided in parentheses. For equal, the test is true if any of the comparisons are equal. For not-equal, the test is true if all of the comparisons are not-equal. For example,

```
LRECL=(80,133)
```

will select datasets whose LRECL is either 80 or 133. Also, a given FIELD name may be specified more than once with several different operators; all of the tests must be true for the dataset to be selected (unless all of the operators are equal (=, .EQ.) when it is sufficient for any one of the tests on that field to be true). For example:

```
SIZE>50,SIZE<100,DSORG=PS,DSORG=PO
```

will select any PS or PO datasets between 50 and 100 tracks in size.

When processing data sets from source selected by DATATYPE= on the PRINT statement, each data set will be passed through this filtering process:

- If any EXCLUDE or XEXCLUDE statements are present, the data set's FIELD values will be compared to each statement. If **all** of the tests on a given statement are true for that data set, it will be excluded from processing.
- If any SELECT or XSELECT statements are present then for any data set which was not excluded, the data set's FIELD values will be compared to each of them. If **all** of the tests on a given statement are true for that data set, it will be processed.
- If there are no SELECT or XSELECT statements present then all data sets which are not excluded will be processed. If none of the above statements are present, then all data sets found on FDREPORT's input will be processed.

The order of the (X)SELECT and (X)EXCLUDE statements is not significant. However, **each statement operates independently**. For example, if you want to select all PDSs on volumes starting with ABC, code:

```
SELECT DSORG=PO,VOL=ABC* <— CORRECT
```

If you code:

```
SELECT DSORG=PO          <— INCORRECT
SELECT VOL=ABC*
```

it will select **all** of the data sets on those volumes, not just the PDSs.

FIELD NAME OPERANDS

In addition to the operands described in this section, you can also include any of the FIELDS described by the tables in [Section 54.31](#) if they are marked as valid for SELECT or XSELECT.

- If an "E" is shown in the SELECT column, then that FIELD may be used on SELECT or EXCLUDE statements and the comparison will be true if the value of the FIELD exactly equals the value you provide.
- If a "G" is shown in the SELECT column, then that FIELD may be used on SELECT or EXCLUDE statements but the comparison will be true if the value of the FIELD is equal to or greater than the value you provide.
- If a "E" is shown in the XSELECT column, then that FIELD may be used on XSELECT or XEXCLUDE statements followed **ONLY** by an equal or not-equal operation.
- If an "A" is shown in the XSELECT column, then that FIELD may be used on XSELECT or XEXCLUDE statements followed by any of the comparison operators supported.

54.13 CONTINUED . . .

For most of the FIELD operands, the values that you specify are obvious. For example, LRECL= takes numeric values specifying logical record lengths. For some FIELDS the rules are not so obvious:

- For FIELDS that represent dates, such as ADATE, EXPDATE and LRDATE, the value must be specified as a Julian date, i.e., year plus day number.
 - For years in the 19xx range you can specify a 2-digit year, e.g., ADATE=95123
 - For any year you can specify a 4-digit year, e.g., EXPDATE=2003123 (required for years beyond 1999)
 - To improve readability, you can insert a period between the year and day, e.g., BKDATE=95.321

The date fields will **not** accept Gregorian dates (month, day, and year), but you can request **display** of dates in Gregorian format with the DATEFORMAT= operand on the PRINT or DEFAULT statement.

- For FIELDS that represent a number of bytes, such as BYTES and PRIBYTES, you may specify values in bytes (a simple number, such as BYTESFRE=5000), kilobytes (a number followed by K, such as BYTES=250K), or megabytes (a number followed by M, such as BYTESUSE=450M). See option HEXBYTES in [Section 54.30](#) for the interpretation of kilobytes and megabytes.
- For FIELDS that represent flags or similar attributes, such as ARCFLAG1, ATTR1, DSIND, and SMSFLAGS, the values that you use on (X)SELECT/(X)EXCLUDE statement may be different from the values that are displayed in a report. The field values are often abbreviated by 1 or 2 characters in a report to save space, but on these statements you may need to put the longer description. For example, ARCFLAG2 will display a value of A if the data set is cataloged for auto-recall, but you would say XSELECT ARCFLAG2=(RECALL) to select on that flag. The tables in [Section 54.31](#) show the report values and the selection values for such fields.

OPERANDS	DSN	Specifies one or more data set names, up to 44 characters each. These names will be compared to those of data sets and VSAM clusters. For the SELECT/EXCLUDE statement only DSN= is supported. Relative GDG generations, e.g., DSN=ABC.GDG(-1) are not supported.
	DSG	Specifies one or more data set name prefixes, up to 44 characters each. The data set names and VSAM cluster names will be checked to see if they begin with the prefix. For the SELECT/EXCLUDE statements, only DSG= is supported. The XDSN= operand, described below, provides a much more flexible way of selecting data set names with a mask. There is a special form of the DSG operand. Leading periods(.) after DSG= or DSG~ indicate that the compare for the prefix name starts after one or more index levels. Each period indicates that one (1) index level is to be bypassed. NOTE: DSN, DSG and XDSN operands can be repeated and/or intermixed.
	MEMNAMES	(XSELECT/XEXCLUDE statements only) For PDSs, matches the data set only if it contains (MEMNAMES=) or does not contain (MEMNAMES~) the members specified. The member names can be specified as masks, as described under the VOL= operand in this section.

54.13 CONTINUED . . .

STORGRP	(XSELECT/XEXCLUDE statements only) Selects (STORGRP=) or excludes (STORGRP¬=) all online volumes defined as belonging to the specified SMS storage group. A storage group prefix can be specified by following the prefix with an asterisk, e.g., STORGRP=DB*. Multiple storage groups and/or prefixes may be specified by enclosing them in parentheses, separated by commas. Valid only on systems with SMS active.
UNIT	(XSELECT/XEXCLUDE statements only) Selects online volumes based on their device address, up to 4 hexadecimal digits. All comparison operands are valid (e.g., UNIT>=140).
UNITNAME	(XSELECT/XEXCLUDE statements only) Selects (UNITNAME=) or excludes (UNITNAME¬=) all online volumes which are mounted on a DASD unit which is included in the specified generic (e.g., UNITNAME=3380) or esoteric name (e.g., UNITNAME=SYSDA). This is limited to names which are valid for UNIT= in JCL at your installation. Multiple units may be specified by enclosing them in parentheses, separated by commas.

54.13 CONTINUED . . .

XDSN

(XSELECT/XEXCLUDE statements only) selects or excludes data sets based on a mask tested against the data set name or VSAM cluster name. XDSN= and XDSN~= are supported. This mask may contain:

- any valid (alphanumeric or national) character representing itself
- / (slash) or % (percent) represents a single valid character
- | (vertical bar) represents a single valid alphabetic character
- + (plus) represents a single valid numeric character
- ? (question) represents a single valid national character (#, \$, or @ in the US)
- * (single asterisk) represents zero or more valid characters within one index level
- ** (double asterisk) represents zero or more valid characters contained in one or more index levels (including their periods)
- . (period) represents a period (index level) in the data set name except for the special cases below
- *. (double asterisk, period) at the beginning of the string represents one or more index levels at the beginning of the dsname
- .** (period, double asterisk) at the end of the string represents one or more index levels at the end of the dsname
- .** (period, double asterisk, period) in the middle of the string represents either a SINGLE period or one or more index levels at that point.

Examples:

XDSN=A.** selects any data set with a first index level of exactly "A".

XDSN=A** selects any data set with a first index level that begins with "A".

XDSN=A*.* is equivalent to XDSN=A**

XDSN=A*.*.* selects any data set with a first index level beginning with "A" which has at least 3 index levels.

XDSN=A*.*.* selects any data set with a first index level beginning with "A" which has exactly 3 index levels.

XDSN=**.*.CNTL selects any data set of at least two index levels ending in ".CNTL".

XDSN=///.* selects any data set starting with a first index level exactly 3 characters long.

Relative GDG generations, e.g., XDSN=*.MASTER(0) are **not** supported.

NOTE: DSN, DSG and XDSN operands can be repeated and/or intermixed.

54.13 CONTINUED . . .

VOL

Specifies 1 or more DASD volume serial numbers. Volume serials may be specified by a mask. This mask may contain:

- any valid (alphanumeric or national) character representing itself
- / (slash) or % (percent) represents a single valid character
- | (vertical bar) represents a single valid alphabetic character
- + (plus) represents a single valid numeric character
- ? (question) represents a single valid national character (#, \$, or @ in the US)
- * (single asterisk) represents zero or more valid characters

Examples: VOL=X**Z VOL=/ /PR* VOL=TSO+++

Note: This masking syntax is also supported for the following data field operands: BKSUFFIX, BKVOL, DATACLAS, DSSN, EXCPEXIT, LMJOB, MGMTCLAS, MEMNAMES, STORCLAS, STORGRP and SYSCODE.

For the SELECT/EXCLUDE statement only VOL= is supported; XSELECT/XEXCLUDE support both VOL= and VOL¬=. Multiple volume serials and/or masks may be specified by enclosing them in parentheses, separated by commas. For example, VOL=(SYS123,TSO+++)

The VOL= operand may be repeated on a single statement; all of the specified volumes will be selected (or excluded). On XSELECT/XEXCLUDE statements, VOL=, UNIT=, UNITNAME= and STORGRP= may be specified separately, repeated, and/or intermixed on one statement to specify the volumes to be processed. If none of these operands are specified, then the volumes on which the data sets reside will not be a criteria for selection.

NOTE: FDREPORT will dynamically allocate any online disk volume identified by VOL, UNIT, UNITNAME and/or STORGRP if a DISKxxxx DD statement is not specified for that volume. The ENABLE=ONLINE operand of the PRINT statement is not required.

54.14 FDREPORT REPORT STATEMENT

REPORT **FIELD=(field1,field2,...,fieldn)**

,RESET|NORESET

REPORT STATEMENT	The REPORT statement specifies the FIELDS which will be included in the output generated by the PRINT statement. The FIELDS will be placed on the report in the order specified. The appropriate heading will be automatically generated for each FIELD selected unless the HEADING statement is present. One blank will be inserted between adjacent fields unless overridden by SPn. If the FIELDS selected generate a line greater than the page width, FDREPORT will truncate the line and issue a warning message. If ENABLE=AUTOSTACK is specified on the PRINT or DEFAULT statement, FDREPORT will stack as many related FIELDS with matching print length and data type as needed to fit the report within the page width. If stacking does not generate a line within the page width, FDREPORT will truncate the line and issue a warning message.
-----------------------------	---

DEFAULT REPORT	<p>If the REPORT statement is not specified, FDREPORT will generate a report with the following FIELDS: SPLDSN, VOL, DSORG, RECFM, BLKSIZE, LRECL, SIZE, SIZEFREE, %FREE.</p> <p>If DATATYPE=VOLDATA, the default report will be: VOL, DEVTYPE, UNIT, VLINDSTA, VLUSEATR, VLTRKVOL, VLALOTRK, VL%UTRKS, VLFREETRK, VLLRGCYL, VLVTOCTR, VLDSCB1 and VL%UDSCB.</p>
---------------------------	--

OPERAND	FIELD=	Specifies one or more names of the FIELDS to be printed. Hundreds of field names are available, in 8 groups: VTOC, VSAM, SMS, PDS, IAM, ABR BACKUP, GENERATED and VOLUME. Fields from any group may be specified. Any which do not apply to a given report line will be blank or zero.
----------------	---------------	--

All of the FIELD names are documented in the tables in [Section 54.31](#) are valid. If SPn is specified as one of the FIELD operands, FDREPORT will insert the number of spaces specified by n (from 0 to 9) between the previous FIELD specified and the next FIELD.

The following example specifies that three (3) blanks be inserted between the FIELD VOL and DATES.

EXAMPLE: REPORT FIELD=(DSN,VOL,SP3,DATES)

The default is one (1) blank between fields.

RESET	RESET (the default) indicates that the list of fields on this REPORT statement completely replaces the values on any preceding REPORT statement.
NORESET	

NORESET indicates that these fields should be added to the list of fields currently in effect from preceding REPORT statements.

54.15 FDREPORT SORT STATEMENT

SORT **FIELD=**(field1,field2,...)
 ,SEQUENCE=(A|D,...)
 ,BREAK=(NO|EJ|RPM|SEJ|SP|SSP|SUB|YES,...)
 ,RESET|NORESET

**SORT
STATEMENT**

The SORT statement specifies the fields to be used to SORT the data selected. You can specify one or more fields to be used for the SORT. Most of the fields which can be reported on can be selected whether or not that field is to be printed.

If this statement is not specified, the SORT=COMBINE and SORT=YES options on PRINT and DEFAULT, and several other reporting options which require sorting, will invoke sorting with default sorting parameters.

Sorting may require that you specify certain DD statements for your SORT product (see JCL REQUIREMENTS in [Section 54.02](#)) or FDREPORT may dynamically allocate all required SORT files (see SORTALLOC= in [Section 54.19](#)).

The BREAK statement, described in the next section, provides a simpler alternative to the SORT statement when all fields are to be sorted in ascending order and a standard control break is to be taken on every field.

OPERANDS

FIELD= Specifies one or more FIELDS to be used by FDREPORT for the SORT. If more than one FIELD is specified, they are sorted in the order specified.

The tables in [Section 54.31](#) document the available FIELD names; check the SORT column to see which are supported for sorting.

SEQUENCE= Specifies the SEQUENCE in which the corresponding SORT FIELD is to be SORTed.

A – ascending sequence.

D – descending sequence.

If more than one FIELD was specified, each value coded is associated with the relative position of the FIELD specified.

For example, if FIELD=(A,B,C) is specified with SEQUENCE=(D,A,D), FIELD A will be SORTed descending, FIELD B ascending and FIELD C descending.

The default is A (ascending).

54.15 CONTINUED . . .

BREAK= Specifies if a change in the value of a SORT FIELD is to cause a control break, and what action to take at that break.

NO – no control break.

EJ – page eject only.

RPM – reset punch mask (reinitialize punch mask processing, [see Section 54.18](#))

SEJ – subtotal summary fields and page eject.

SP – space one line only.

SSP – subtotal summary fields and space one line.

SUB – subtotal summary fields.

YES – subtotal summary fields, space one line, and reset punch mask.

If more than one FIELD was specified, each value coded is associated with the relative position of the FIELD specified.

For example, if FIELD=(A,B,C) is specified with BREAK=(YES,NO,SSP), a change in FIELD A will do summary processing, FIELD B will not, and FIELD C will cause subtotals with a blank line.

The default is NO for all FIELDS.

RESET **NORESET** **RESET** (the default) indicates that the list of fields on this SORT statement completely replaces the values on any preceding SORT or BREAK statement.

NORESET indicates that these fields should be added to the list of fields currently in effect from preceding SORT or BREAK statements.

54.16 FDREPORT BREAK STATEMENT

BREAK **FIELD=**(field1,field2,...)

,RESET|NORESET

BREAK STATEMENT The BREAK statement provides a simple alternative to the SORT statement when you want to sort every field in ascending order and take a standard control break when any value changes. In other words,

BREAK FIELD=(A,B,C)

is equivalent to

SORT FIELD=(A,B,C),BREAK=(YES,YES,YES)

Details are found in [Section 54.15](#) (SORT statement)

OPERANDS **FIELD=** Specifies one or more FIELDS to be used by FDREPORT for the SORT. If more than one FIELD is specified, they are sorted in the order specified. All fields are sorted in ascending order.

The tables in [Section 54.31](#) document the available FIELD names; check the SORT column to see which are supported for sorting.

RESET **RESET** (the default) indicates that the list of fields on this BREAK statement completely replaces the values on any preceding BREAK or SORT statement.

NORESET

NORESET indicates that these fields should be added to the list of fields currently in effect from preceding BREAK or SORT statements.

54.17 FDREPORT SUMMARY STATEMENT

SUMMARY **FIELD=**(field1,field2,...)

,RESET|NORESET

SUMMARY STATEMENT The SUMMARY statement provides summary reports based on the values of various fields reported by FDREPORT. Summary reports will be printed on the ABRSUM DD if present, or will be interspersed with the generated report on ABRMAP or SYSPRINT if ABRSUM is absent.

There are 3 kinds of summarization done, depending on the characteristics of the field being summarized. Fields which always have a unique value (such as DSN) will simply print all unique values encountered. Certain numeric fields (such as SIZEFREE) will be totaled. Other character and numeric fields (such as BLKSIZE) will report the number of unique values encountered and the count of times that value was found.

Summaries will be printed at control breaks (see the BREAK= operand of the SORT statement) and at the end of the report. The final summary will contain totals for the entire report; control break summaries will contain values since the last control break.

Examples of summaries are shown in the FDREPORT examples.

OPERANDS **FIELD=** Specifies one or more FIELDS to be summarized.

The tables in [Section 54.31](#) document the available FIELD names; check the SUMMARY column to see which are supported for summarization and the type of summary which will be done.

RESET **RESET** (the default) indicates that the list of fields on this SUMMARY statement completely replaces the values on any preceding SUMMARY statement.

NORESET **NORESET** indicates that these fields should be added to the list of fields currently in effect from preceding SUMMARY statements.

54.18 FDREPORT PUNCH STATEMENT

PUNCH	ECHO
	, FDRLIB =ddname
	, FIELDPREFIX =c
	, FIELDSUFFIX =c
	, MASKNAME =member
	, MAXSTATEMENTS =nnnn
	, SYMBOLS =(symbol,symbol,...)
	, VALUES =(value,value,...)

PUNCH STATEMENT The PUNCH statement specifies the location of a mask used to define the "punch" output generated by FDREPORT for the RPTYPE=SELPCH report and is ignored if RPTYPE=SELPCH is not specified.

The punch mask input must be a data set that has the attributes: RECFM=FB,LRECL=80. It may be either sequential (DSORG=PS) or partitioned (DSORG=PO). The PUNCH statement uses the presence or absence of the MASKNAME= operand to differentiate between sequential or partitioned data sets. If MASKNAME= is specified, the data set is assumed to be partitioned and MASKNAME= provides the member name. If MASKNAME= is omitted, the data set is assumed to be sequential (the FDRLIB= operand must be provided to specify the input DD name). The punch mask may also be an input stream (DD *) data set.

For each record processed by FDREPORT, it will scan the punch mask, copying each mask record to the punch output DD statement (SYSPUNCH), but if any FDREPORT field names or special names are found in the mask, the current value of the name is substituted. These field names and special names must be surrounded by the delimiters specified by the FIELDPREFIX= and FIELDSUFFIX= operands; the defaults for these are < and >, so the field names are usually specified like: <LRECL>. The field names that may appear in the PUNCH mask are found in the tables in [Section 54.31](#) (check the PUNCH column). In addition to the those field names, the following special names may be used:

- \$\$CNTFDR – record counter with one (1) byte significance ranging from 0-9, A-Z.
- \$\$CNT2 – record counter with two (2) byte significance.
- \$\$CNT3 – record counter with three (3) byte significance.
- \$\$CNT4 – record counter with four (4) byte significance.
- \$\$CNT5 – record counter with five (5) byte significance.
- \$\$CNT6 – record counter with six (6) byte significance.
- \$\$LPCNT2 – loop counter with two (2) byte significance.
- \$\$LPCNT3 – loop counter with three (3) byte significance.
- \$\$PRCNT2 – prior record counter with two (2) byte significance.
- \$\$PRCNT3 – prior record counter with three (3) byte significance.
- \$\$+CNT2 – prior loop counter with two (2) byte significance.
- \$\$+CNT3 – prior loop counter with three (3) byte significance.
- TODAY – current date in Julian format, i.e., yyddd.
- user-specified variable names (the SYMBOLS= operand)

The counters can be used to generate names (such as step names, DD names) which are unique. The \$\$CNTxxx counters are incremented for each FDREPORT record (data set name) passed through the punch mask. The \$\$LPCNTx counters increment each time the punch mask is initialized (such as a control break). The \$\$Pxxxx counters contain the prior value of the associated counter (for example, for generating refer-back JCL).

54.18 CONTINUED . . .

There are five (5) positional operands that can be used within the mask to permit one time generation of a segment of the mask. The operands must begin in column one (1) and are described as follows:

)REPRO – statements following are processed only once.

)PREFIX – statements following are processed once per loop, at the beginning of the loop, controlled by BREAK=RPM/YES on SORT statement).

)ENDPREFIX – terminates the loop prefix statements.

)SUFFIX – statements following are processed once per loop, at the end of the loop.

)INCREMENTLOOP – increment current loop count. Only executed once.

An example of a punch mask:

```
)PREFIX
//USER1LST      JOB   (USER1,123),LISTCAT,MSGCLASS=X,
//              CLASS=C
)ENDPREFIX
//STEP<$$CNT3> EXEC PGM=IDCAMS
//SYSPRINT      DD    SYSOUT=*
//SYSIN         DD    *
LISTCAT ENT(<NAME>) ALL
)SUFFIX
//STEPEND       EXEC  PGM=SOMEPRGM
//SYSPRINT      DD    SYSOUT=*
```

If RPTYPE=SELPCH is specified and no PUNCH statement is present, the default punch mask is:

```
SELECT VOL=<VOL>,DSN=<NAME>
which generates statements acceptable to ABR.
```

OPERANDS	ECHO	Specifies that the punch mask is to be printed on SYSPRINT. The default is the mask will not be printed.
	FDRLIB=	Specifies the ddname to be used when reading punch masks. The default is FDRLIB.
	FIELDPREFIX=	Specifies the character that marks the beginning of a field name in a punch mask. Valid characters are < &\$*%_>:#@ but the default of < should not be overridden unless you need to use that character as data in the mask.
	FIELDSUFFIX=	Specifies the character that marks the end of a field name in a punch mask. Valid characters are < &\$*%_>:#@ plus a blank but the default of > should not be overridden unless you need to use that character as data in the mask.
	MASKNAME=	Specifies the member of the data set pointed to by the ddname specified in the operand FDRLIB is to be used as a mask for the RPTYPE=SELPCH output. The FDRLIB data set must be partitioned (PO). Either FDRLIB= or MASKNAME= must be specified.

54.18 CONTINUED . . .

- MAXSTATEMENTS=** Specifies the number of control statements that can be present in any single punch mask. The number may be a value from 10 to 4000, inclusive.
- The default is 400 statements.
- SYMBOLS=** Specifies 1 or more user-defined field names which can be specified in the punch mask. The corresponding value of each field name must be specified by a **VALUES=** operand. This can be used to insert variable information (e.g., security or accounting parameters) in a punch mask.
- VALUES=** Used with the **SYMBOLS=** operand to specify the values of the user-defined fields to be substituted in the punch mask. There must be as many values specified as there are field names in the **SYMBOLS=** operand.

54.19 FDREPORT PRINT STATEMENT

PRINT	ALIASLEVEL=n ,ALTDATADDNAME=ddname ,ARCLIMIT=nnnn ,BYTEFORMAT=BYTES <u>KILOBYTES</u> MEGABYTES ,COPY=BOTH <u>EITHER</u> 1 2 ,DATADDNAME=ddname ,DATATYPE=ARCHIVE BCDS CATALOG CATARCH CATVTOC EXTRACT MCDS TVTOC VOLDATA <u>VTOC</u> ,DATEFORMAT=YYDDD MMDDYY DDMMYY MMDDYYYY DDMMYYYY YYYYDDD ,DISABLE=(option1,option2,...) ,ENABLE=(option1,option2,...) ,DSKIP=n ,EXTRACTDSN=dsname ,EXTRACTMEMBER=membername ,FORMAT=CRT PRT ,INDEXNUM=nn ,LBPZERO= <u>INVALID</u> VALID ,LINECNT=nn	,MAXGDG=nnn ,MAXONLINE=nnnn ,OLDBACKUP=(nn,nn,...) ALL <u>CUR</u> ,PAGEWIDTH=nnn PCHDATEFORMAT= <u>YYDDD</u> MMDDYY DDMMYY MMDDYYYY DDMMYYYY YYYYDDD ,RPTYPE=ABRVTOC ARCHIVE COMPDATA DATA <u>GENERATE</u> HEX NONE OSVTOC SELPCH TABLE XREF ,SKIP=n ,SORT=COMBINE <u>NO</u> YES ,SORTALLOC=(option1,option2,...) ,SORTCORE=nnnnnnn ,SORTLIB=dsname ,SORTMSGDDNAME=ddname ,SORTPFX=cccc ,SUM=YES <u>NO</u> INDEX ,SUMDEVICE= <u>BASE</u> UNIQUE ,SYSUTSPACE=nnnn ,SYSUTUNIT=unitname ,TITLE=LEFT <u>CENTER</u> RIGHT ,WORKDDNAMES=n ,WORKSPACE=nnnn ,WORKUNIT=unitname
-------	--	---

PRINT STATEMENT The PRINT statement actually generates the report defined by previous FDREPORT statements and operands on PRINT itself. FDREPORT will perform the PRINT using the characteristics specified on the preceding SELECT, REPORT, DEFAULT and/or SORT statements. A PRINT statement must be specified; multiple PRINT statements may be given to produce various reports in one execution of FDREPORT. All parameters which define the report to be generated (TITLE, HEADING, REPORT, SORT, SUMMARY, (X)SELECT, (X)EXCLUDE, and/or PUNCH) **must precede** the PRINT statement.

54.19 CONTINUED . . .

OPERANDS	ALIASLEVEL=	<p>Specifies an alias level, from 0 to 4, to be used when selecting data sets from the system catalogs. If 0, it uses the Multi-Level Alias (MLA) level currently in effect in the Catalog Address Space on your system. If 1 through 4 is specified, it will act as though that were the current MLA level active on your system; this provides a way to test an MLA catalog structure before you activate it.</p> <p>Default is 0.</p>
	ALTDATADDNAME=	<p>Specifies the DD name of the file to be used for the input of FDREPORT internal records when RPTYPE=DATA (or COMPDATA) and DATATYPE=EXTRACT are both specified.</p> <p>Default is SYSUT1.</p>
	ARCLIMIT=	<p>When DATATYPE=ARCHIVE or CATARCH is specified, only data sets which were archived within the last "nnnn" days will be selected.</p> <p>Default is 0, which disables ARCLIMIT checking.</p>
	BYTEFORMAT=	<p>Specifies the units in which fields which represent a number of bytes are to be processed:</p> <p>BYTES – actual bytes.</p> <p>KILOBYTES – units of kilobytes (1000 bytes).</p> <p>MEGABYTES – units of megabytes (1,000,000 bytes).</p> <p>The fields affected include: BYTES, PRIBYTES, SECBYTES, BYTESUSE, BYTESFRE (see Tables 1-8 in Section 54.31 for details). Kilo and megabytes will be in units of 1024 and 1048576 if ENABLE=HEXBYTES is in effect.</p> <p>For printing (REPORT statement), this operand affects both the units in which the field is printed and the width of the field in the report (10 columns for BYTES, 7 for KILOBYTES, 4 for MEGABYTES).</p> <p>The default is KILOBYTES.</p>
	COPY=	<p>When fields relating to ABR backup or archive information are included in the report, specifies which ABR copy (1 or 2) will appear in the report. Values are:</p> <p>1 – information about COPY 1 will be used.</p> <p>2 – information about COPY 2 will be used.</p> <p>EITHER – COPY 1 information will be used, if a COPY 1 backup is recorded, otherwise COPY 2 will be used.</p> <p>BOTH – both copies will be used (if both are cataloged). Will print 2 lines in the report.</p> <p>Default is COPY=EITHER.</p>

54.19 CONTINUED . . .

DATADDNAME= Specifies the DD name of the file to be used for the output of FDREPORT internal records when RPTYPE=DATA is specified, or for the input of those records when DATATYPE=EXTRACT is specified. Default is SYSUT2.

DATATYPE= Specifies the source of FDREPORT's input data. Values are:

- ARCHIVE** – read records from an archive control file. If an ARCHIVE DD statement is present, that archive control file will be read; otherwise the control file name specified in the ABR option table or by ARCDN= will be dynamically allocated.
- BCDS** – report on data extract from a DFSMSHsm or DFHSM BCDS (Backup Control Data Set). [See Section 54.32](#) for details.
- CATALOG** – data set records will be extracted from system catalogs.
- CATARCH** – data set records will be extracted from system catalogs; only those data sets cataloged for auto-recall will be selected. The extracted names will be used to select data sets from the Archive Control File (see ARCHIVE above).
- CATVTOC** – data set records will be extracted from system catalogs. The extracted names will be used to select data sets from the VTOCs of the volumes extracted from the catalog (see VTOC below).
- EXTRACT** – FDREPORT internal records will be read from the SYSUT2 DD statement. This data set must have been created by FDREPORT using RPTYPE=DATA or RPTYPE=COMPDATA in this step or a previous step or job. This may be used to produce several reports from the same set of data without the overhead of reconstructing that data. For extract files created with RPTYPE=DATA only, SYSUT2 may be a concatenated DD statement, allowing you to read several extract files as one.
- MCDS** – report on data extract from a DFSMSHsm or DFHSM MCDS (Migration Control Data Set). [See Section 54.32](#) for details.
- TVTOC** – information is extracted from a backup file created by FDR, FDRDSF, FDRABR, or SAR.
- VOLDATA** – volume summary data is gathered for selected volumes using LSPACE, VTOC, VTOCIX, and VVDS, creating a volume record.
- VTOC** – data set information will be read from the VTOCs and VVDSs of volumes selected.

Default is VTOC.

54.19 CONTINUED . . .

DATEFORMAT=	<p>Specifies the format that dates will be printed in generated reports. Valid values are:</p> <p>YYDDD – Julian format ("yy.ddd").</p> <p>YYYYDDD – Julian format with 4-digit year ("yyyy.ddd").</p> <p>MMDDYY – Gregorian format, month first ("mm/dd/yy").</p> <p>MMDDYYYY – Gregorian format, month first with 4-digit year ("mm/dd/yyyy").</p> <p>DDMMYY – Gregorian format, day first ("dd/mm/yy").</p> <p>DDMMYYYY – Gregorian format, day first with 4-digit year ("dd/mm/yyyy").</p> <p>If one of the Gregorian formats is chosen, then date fields will be 8 columns long in reports, instead of the 6 columns shown in the tables which follow. If a 4-digit year format is chosen, the date fields are an additional 2 bytes. This affects only printing; sorting is always in Julian ("yyyddd") format and selection operands can use Julian in either 2- or 4-digit year format.</p> <p>Default is YYYYDDD (Julian format).</p> <p>Note: Prior to V5.3, the default was DATEFORMAT=YYDDD (2-digit year); in V5.3 and above, the default is DATEFORMAT=YYYYDDD (4-digit year). This may affect the format of existing FDREPORT jobstreams if they depend on the length of the default date format. See Section 54.10 (the DEFAULT statement) for instructions on permanently changing the default.</p>
DISABLE=	<p>The FDREPORT options specified are to be disabled for this run. Multiple options can be specified by enclosing the list in parentheses, separated by commas. Available options are described in Section 54.30.</p>
DSKIP=	<p>Specifies the number of blank lines to be inserted between lines representing separate data sets (differs from SKIP= when multiple lines are required for one data set).</p> <p>Default is 0.</p>
ENABLE=	<p>The FDREPORT options specified are to be enabled for this run. Multiple options can be specified by enclosing the list in parentheses, separated by commas. Available options are described in Section 54.30.</p>
EXTRACTDSNAME=	<p>Specifies the data set name of an existing data set to be used for FDREPORT extract data input (if DATATYPE=EXTRACT) or output (RPTYPE=DATA or COMPDATA). This data set will be dynamically allocated as DD name SYSUT2 (or whatever is specified by DATADDNAME=).</p>
EXTRACTMEMBER=	<p>Used in conjunction with EXTRACTDSNAME=. If the data set pointed to by EXTRACTDSNAME= is a PDS, this specifies a member name in that PDS to be used for input or output.</p>

54.19 CONTINUED . . .

FORMAT=	<p>Specifies that the report is to be prepared using other than the default format selected by the program.</p> <p>PRT – defaults to PAGEWIDTH=120 to generate a report formatted for printing.</p> <p>CRT – defaults to PAGEWIDTH=78 to generate a report formatted for viewing on a terminal. The line length can be overridden by the PAGEWIDTH= operand.</p> <p>Default is CRT if executed under TSO, or PRT otherwise.</p>
INDEXNUM=	<p>Field name INDEX (see Table 5 in Section 54.31) allows sorting, summary, etc., on an index level extracted from each data set name. INDEXNUM= specifies which index level is to be extracted (1 is first, 2 is second, etc.).</p> <p>The default is 1, the high-level index.</p>
LBPZERO=	<p>Specifies how a PS (sequential) data set whose last block pointer is all zeros is to be treated for used and free track calculations.</p> <p>VALID means that such data sets are considered to have no tracks used and all tracks free.</p> <p>INVALID means they are considered to be entirely used and no tracks free.</p> <p>Since most access methods maintain a valid last block pointer in the Format 1 DSCB, LBPZERO=VALID will produce correct reports for almost all data sets so it is recommended.</p> <p>Default is INVALID.</p>
LINECNT=	<p>Specifies the maximum number of lines each report page can contain, from 28 to 99 inclusive.</p> <p>The default is 58 .</p>
MAXGDG=	<p>When reading from the system catalogs (DATATYPE=CATALOG, CATVTOC, or CATARCH), specifies the maximum number of generations of each GDG (Generation Data Group) that will be selected. Only the most recent "nnn" generations will be processed.</p> <p>The default is that all generations will be processed.</p>
MAXONLINE=	<p>Specifies the maximum number of disk volumes which can be processed by this PRINT statement. The value can be from 5 to 9000. If the number of disk volumes selected exceeds this number, an error will occur. This is a special concern when using ENABLE=ONLINE in installations with large DASD installations.</p> <p>The default is 256 disk volumes. However, if ENABLE=AUTOMAXON is in effect, MAXONLINE= will be automatically set to the number of disk volumes online to this system.</p>

54.19 CONTINUED . . .

OLDBACKUP=

when ABR backup information is requested, and old backups are being recorded, specifies which old backup information is to be printed.

Values are:

(nn,...,nn) – requests specific old backups. "nn" are numbers from 00 to 13 (00 being the most recent backup, 13 being the oldest).

ALL – requests all old backups existing for a data set are to be printed.

CUR – requests that only the current backup is to be printed.

If multiple backups exist and are requested for a data set, the report will contain multiple lines for the data sets.

Default is CUR.

PAGEWIDTH=

Specifies the number of print positions (from 50 to 200) to be used in creating the report, excluding the printer control character in column 1. Heading lines and formatted data fields should fit within this limit. If the data fields requested exceed the PAGEWIDTH, FDREPORT will print all of the data which fits and display a warning message.

Default is 78 if FORMAT=CRT and 120 if FORMAT=PRT.

NOTE: PAGEWIDTH is ignored if the report is printed on SYSPRINT. ABRMAP must be present for PAGEWIDTH to be honored.

PCHDATEFORMAT=

Specifies the format that dates will be displayed in "punched" output generated by RPTYPE=SELPCH. Valid values are:

YYDDD -- Julian format ("yy/ddd").

YYYYDDD -- Julian format with 4-digit year ("yyyy/ddd").

MMDDYY -- Gregorian format, month first ("mm/dd/yy").

MMDDYYYY -- Gregorian format, month first with 4-digit year ("mm/dd/yyyy").

DDMMYY -- Gregorian format, day first ("dd/mm/yy").

DDMMYYYY -- Gregorian format, day first with 4-digit year ("dd/mm/yyyy").

54.19 CONTINUED . . .

RPTYPE=

Specifies the report type to be generated by FDREPORT, and can request certain predefined report formats or special processing. When predefined reports are requested, the REPORT and HEADING statements and the AUTOSTACK operand are ignored. Values are:

ABRVTOC – a report in standard ABR VTOC format (equivalent to PRINT VTOC) is printed.

ARCHIVE – a standard ABR ARCHIVE report (equivalent to PRINT ARCHIVE) is printed; intended for use with DATATYPE=ARCHIVE.

DATA – no report will be printed, but FDREPORT internal records for every data set or volume selected will be written to the extract data file on DD name SYSUT2. This data may be read as input to FDREPORT in this or another step (DATATYPE=EXTRACT) allowing multiple reports to be generated without the overhead of reconstructing the data. With RPTYPE=DATA, FDREPORT will include in the extract file all relevant data fields from its primary source specified by DATATYPE=; for example, if you specify DATATYPE=VTOC or CATVTOC, all fields generated from the VTOC and VVDS are available. If you need additional fields from another source, include a REPORT statement requesting at least one such field, e.g.,

```
REPORT FIELD=(CATALOG,DIRBLOCKS)
```

which will gather ALL catalog and directory fields in addition to the VTOC/VVDS fields.

COMPDATA – similar to RPTYPE=DATA, except that only the fields named in a REPORT statement will be written to the extract file; this makes the extract file much smaller. Note that an extract file created by DATATYPE=COMPDATA cannot be part of a concatenated DD statement when used as input.

GENERATE – generates a customized report based on FDREPORT options (default).

HEX – prints records from the ARCHIVE control file (if DATATYPE=ARCHIVE) or records generated by FDREPORT (for all other data types) in a dump (hex and character) format.

NONE – suppresses the detail report, allowing only summaries to be printed (SUM=YES, SUM=INDEX or the SUMMARY statement should also be specified).

OSVTOC – a report containing information equivalent to an IEHLIST LISTVTOC statement (but not in the same format) is printed.

SELPCH – no report will be printed, but for each data set selected a statement will be written to SYSPUNCH in the format specified by the MASK operand of the PUNCH statement, or, by default, in the format:

```
SELECT VOL=volser,DSN=dsname
```

The SYSPUNCH data set may be passed to a following step, possibly a ABR step, allowing the enhanced selection facilities of FDREPORT to select data sets to be processed by ABR or other programs, or it may be submitted directly to a JES internal reader if a complete jobstream is generated by the punch mask. A SORT will be forced to properly handle multi-volume data sets unless SORT=NO is specified.

54.19 CONTINUED . . .

TABLE – generates output designed to be read by other programs. No titles will be printed and the report is not limited by the PAGEWIDTH operand. One set of column headings are printed unless DISABLE=HEADINGS is specified. Records up to 32760 bytes long may be created (depending on the characteristics of the ABRMAP data set). The ABRMAP DD is required, this report cannot go to SYSPRINT. Fields requested by the REPORT statement will appear in sequence with single columns between them.

a report in standard ABR VTOC/BACKUP XREF format (equivalent to PRINT BACKUP) is printed.

SKIP=

Specifies the number of blank lines to be inserted between report lines. The number can be a value from 0 (single space) to 3.

The default is 0.

SORT=

Specifies if the data is to be SORTed. However, sorting will be forced by a SORT statement and by some other options which require sorting.

YES – sorts by disk volume serial. If a SORT statement is present, and the first sort field is not VOL, VOL will be temporarily inserted as the first field. If no SORT is present, data is sorted by volume serial and data set name, both ascending.

NO – no sorting is performed. SORT statements and implied sorts are ignored.

COMBINE – if no SORT statement is present, sorts by data set name, volume sequence number, and volume serial number, all ascending, to produce a report sorted by data set name across all volumes. Ignored if a SORT statement is present.

If sorting is required, any DD statements required by your installation's SORT product must be included in the FDREPORT step unless the SORTALLOC operand is specified.

Default is NO unless a SORT statement precedes the PRINT statement or SUM or RPTYPE implies sorting.

54.19 CONTINUED . . .

SORTALLOC=

If sorting is specified or forced, this specifies if FDREPORT will dynamically allocate some or all files required by your system sort product.

NO – do not dynamically allocate SORT related data sets. If sorting is required, any necessary DD statements must be included in the JCL.

SORTLIB – dynamically allocate the SORTLIB data set using the value in the operand SORTLIB for the dsname. Users of the SYNC SORT product should see the note under the SORTLIB= operand.

SORTMSG – dynamically allocate the SORT message output to SYSOUT using the value in the operand SORTMSGDDNAME for the ddname.

SORTWORK – dynamically allocate the number of SORTWKnn data sets specified in the operand WORKDDNAMES, using the value in the operand WORKUNIT as the unitname.

YES – dynamically allocate all of the above.

CYL – if SORTWKxx files are allocated, the allocation will be in cylinders.

TRK – if SORTWKxx files are allocated, the allocation will be in tracks. More than one option can be specified, in parentheses, separated by commas, e.g., SORTALLOC=(YES,CYL).

Default is (NO,TRK).

SORTCORE=

Specifies the amount of storage the program SORT is to use if external sorting is required. The number may be from 10000 to 8000000 inclusive.

The default is taken from the FDR Global Option Table and is usually 100000.

SORTLIB=

Specifies the dataset name to be allocated to the DDNAME SORTLIB. The default is 'SYS1.SORTLIB'.

NOTE for SYNC SORT users: The SYNC SORT product does not require a SORTLIB, so SYS1.SORTLIB may not exist on your system. If SORTALLOC=ALL is specified, FDREPORT will attempt to allocate it, and it may fail. To circumvent this, you can either create an empty PDS called SYS1.SORTLIB, or override the SORTLIB= operand to specify some other PDS.

SORTMSGDDNAME=

Specifies the DDNAME to be used by the program SORT if messages are to be printed.

The default is SYSOUT.

SORTPFX=

Specifies the DDNAME prefix to be used by the program SORT if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.

The default is taken from the FDR Global Option Table and is usually "SORT".

54.19 CONTINUED . . .

SUM=	<p>controls the printing of summary reports, in conjunction with the SUMMARY statement. Summaries will be printed at indicated points in the detail report (on ABRMAP or SYSPRINT) unless an ABRSUM DD Statement is present, when the summary report will be printed on ABRSUM.</p> <p>NO – no summary reports are printed.</p> <p>YES – summary reports are printed as specified by the SUMMARY statement. If no SUMMARY is present, this is assumed:</p> <p>SUMMARY FIELD=(DSN,NOEXTENT,SIZE,SIZEFREE,SIZEUSED)</p> <p>INDEX – summary reports are printed for each high-level data set index encountered. The SUMMARY statement is honored, or if absent defaults as shown above. SUM=INDEX forces a sort on data set name unless SORT=NO is specified.</p> <p>Default is NO unless a SUMMARY statement precedes the PRINT statement.</p>
SUMDEVICE=	<p>specifies how summaries by device type are to be handled if there is more than one density for the device.</p> <p>BASE – summary is at the base level, e.g., 3390-1, 3390-2, 3390-3 and 3390-9 are summarized as 3390.</p> <p>UNIQUE – summary is by individual device type, e.g., 3390-1, 3390-2, etc. will be summarized separately.</p> <p>Default is BASE.</p>
SYSUTSPACE=	<p>Specifies the number of tracks or cylinders (controlled by the TRK and CYL operands of SORTALLOC=) to allocate when the temporary SYSUT2 extract file are dynamically allocated, from 10 to 1000.</p> <p>The default is 100.</p>
WORKUNIT=	<p>Specifies the unit name (1 to 8 characters) to use when dynamically allocating the temporary SYSUT2 extract file. It must be a value valid for UNIT= in JCL, and the volumes included on those units must include some in STORAGE or PUBLIC status for the allocation to be successful.</p> <p>The default is SYSALLDA which is valid on all MVS systems and includes all DASD devices.</p>
TITLE=	<p>Controls the placement of the title line within the current pagewidth. Values are:</p> <p>LEFT – left aligned.</p> <p>CENTER – centered within the pagewidth.</p> <p>RIGHT – right aligned.</p> <p>Default is CENTER.</p>

54.19 CONTINUED . . .

WORKDDNAMES=	<p>Specifies the number of SORT work DDnames to allocate, from 1 to 5, inclusive.</p> <p>The default is 3.</p>
WORKSPACE=	<p>Specifies the number of tracks or cylinders (controlled by the TRK and CYL operands of SORTALLOC=) to allocate to each of the SORT work files, from 10 to 1000.</p> <p>The default is 100.</p>
WORKUNIT=	<p>Specifies the unit name (1 to 8 characters) to use when dynamically allocating sort work files if requested by SORTALLOC=. It must be a value valid for UNIT= in JCL, and the volumes included on those units must include some in STORAGE or PUBLIC status for the allocation to be successful.</p> <p>The default is SYSALLDA which is valid on all MVS systems and includes all DASD devices.</p>

54.20 FDREPORT CANCEL STATEMENT

CANCEL	EXCLUDE	,SELECT
	,HEADING	,SORT
	,REPORT	,SUMMARY
	,PUNCH	,TITLE

CANCEL STATEMENT The CANCEL statement negates the effects of all or some prior statement except DEFAULT. This statement is very handy if you have changed your mind about the selection criteria, sort fields, etc. or if you wish to generate a totally different report in the same FDREPORT execution.

If no operands are specified, CANCEL will cancel the effect of all of the prior statements except DEFAULT.

OPERANDS	EXCLUDE	Cancel the current exclusion criteria table as created by the EXCLUDE and XEXCLUDE statements.
	HEADING	Cancel the current HEADING line(s).
	REPORT	Cancel the current REPORT field table.
	PUNCH	Cancel the current PUNCH mask.
	SELECT	Cancel the current selection criteria table as created by the SELECT and XSELECT statements.
	SORT	Cancel the current SORT field table.
	SUMMARY	Cancel the current SUMMARY options.
	TITLE	Cancel the current TITLE line.

54.21 FDREPORT EXECUTE STATEMENT

EXECUTE ECHO

,FDRLIB=ddname

,REPORT=reportname

**EXECUTE
STATEMENT**

The EXECUTE statement reads pre-established FDREPORT report statements from a control statement library. You can setup canned report specifications which any user can execute.

The control statements read by EXECUTE can also be combined with statements in the input stream. For example, you might have SELECT statements to select the data sets to be reported, followed an EXECUTE to read and execute the REPORT and PRINT statements defining the report.

The control statement library must be a data set that has the attributes: RECFM=FB,LRECL=80. It may be either sequential (DSORG=PS) or partitioned (DSORG=PO). The EXECUTE statement uses the presence or absence of the REPORT= operand to differentiate between sequential or partitioned data sets. If REPORT= is specified, the data set is assumed to be partitioned and REPORT= provides the member name. If REPORT= is omitted, the data set is assumed to be sequential.

NOTE: The ISPF SRS dialog, documented starting in [Section 54.50](#), is another way to save selection and/or formatting parameters in a library which can then be executed under TSO or submitted for batch execution.

OPERANDS

ECHO Specifies that the control statements read are to be printed on SYSPRINT.

The default is the statements will not be printed.

FDRLIB= Specifies the DD name FDREPORT is to scan for the control statements to be processed.

If REPORT is specified this DD must point to a partitioned data set. If REPORT is not specified FDREPORT will read this data set sequentially.

The format of this data set must be fixed blocked with an LRECL of 80. The blocksize is set at user discretion.

The default is FDRLIB.

REPORT= Specifies the name of the report. This name must be a member in the library specified by FDRLIB. This member must contain the control statements defining this report in 80 byte images. The user can specify any of the FDREPORT statements in this member except for EXECUTE. If the last statement in the member is PRINT, the user need not specify any other control statements after this Statement.

If REPORT is not specified, FDREPORT will assume that the data set pointed by FDRLIB is a sequential data set.

A 1 to 8 character name may be specified.

54.30 FDREPORT EXECUTION OPTIONS

These are suboperands of the ENABLE= and DISABLE= operands on the PRINT and DEFAULT statements. They control various execution options of FDREPORT and invoke special functions. Because of space limitations, some options which are for special purposes and are of limited use are not documented here. To see the complete list of options, execute FDREPORT with the control statement:

```
HELP COMMAND(PRINT)
```

The defaults shown for each option are the defaults distributed with FDREPORT. However, all of these defaults can be permanently changed as described under the DEFAULT statement in [Section 54.10](#).

AIXCLUSTER	for VSAM alternate indexes (AIXs) the name of the AIX cluster will be displayed as the cluster name instead of the base cluster name. Disabled by default.
ALLFILTER	when scanning the catalog (DATATYPE=CATALOG or CATVTOC or CATARCH), select data sets even if they are cataloged to tape or marked for ABR auto-recall. If disabled, only data sets cataloged to disk without auto-recall are selected. Disabled by default.
AUTOFATDISK	automatically determine if support for IBM 3390-9 DASD and other disks with more than 99999 tracks should be enabled; if any disk with over 99999 tracks is online, the FATDISK option is automatically enabled. Enabled by default.
AUTOMAXON	automatically determine the number of DASD volumes online to this processor and set the MAXONLINE= operand to default to that value. This enables you to report on any number of DASD volumes up to your entire DASD installation. Disabled by default.
AUTOSTACK	if the fields to be printed exceed the page width, fields with similar attributes will be printed stacked one above the other. Disabled by default.
AVERAGING	if summarization is requested, print the average value of all numeric summarized fields. Disabled by default.
CHAREXPDATES	If enabled, expiration dates of 1999.365 through 1999.999 are displayed as "NEVER". An expiration date of 1999.000 (99000) will be displayed as "CATCTL". Disabled by default.
COMPSELECT	if data set name selection does not select an ICF VSAM cluster based on the cluster name, the component names will also be checked; if any match the cluster is selected. Enabled by default.
DAFREESPACE	calculate free space in direct access (DSORG=DA) data sets. Disabled by default.
DATELOCATE	catalog LOCATEs will be issued to get creation/expiration dates for ICF VSAM clusters. Disabled by default.
DATESTAMP	report heading is to contain a date. Enabled by default, except disabled by default under TSO.

54.30 CONTINUED . . .

DIAGNOSEVVDS	when executing any FDRREPORT function that accesses disk VTOCs, for volumes which have a VVDS, FDRREPORT will check the VVDS for duplicate (more than one for same name) or orphan (not in VTOC) VVRs and NVRs. It will also display the IDCAMS statements necessary to correct each error; if the PCHDIAGNOS option (see below) is also enabled, it will write them to the PCHDIAG DD statement as well. Disabled by default.
DIRBLOCKS	causes FDRREPORT to read the directory of any selected PDS, even if directory-related fields have not been requested. Disabled by default, but assumed if directory fields are requested.
DISPLAY	can only be specified on a DEFAULT statement, and causes it to display all current FDRREPORT operand and option defaults. Disabled by default.
DSCBDISPLAY	displays the DSCBs read from VTOCs for all selected data sets. Disabled by default.
DUPDSNCHECK	during PUNCH processing, bypasses duplicate data set names, generating only one PUNCH output for each name; for multi-volume data sets and ICF VSAM clusters, ensures that only one output is generated per data set. It is effective only if the data is sorted by data set name. Enabled by default.
FATDISK	required for proper reporting when processing the IBM 3390-9 DASD or any other disk with more than 65535 tracks. This option will change the size of some of the fields to accommodate the larger disks. Disabled by default, but may be automatically enabled by the AUTOFATDISK option described above.
FORMAT4DSCB	request that the VTOC itself be identified by the name "FORMAT4.DSCB" in reports generated from the VTOC, rather than the default of "****VTOC". Disabled by default.
FOURBYTEUNIT	requests that all device unit addresses be processed as 4-digit values as parameter input and in reports. Disabled by default, but enabled automatically if your system is running a level of MVS that supports 4-digit device addresses. This allows you to convert your reports to the new format even before that support is installed.
GDGONLY	when scanning the catalog (DATATYPE=CATALOG or CATVTOC or CATARCH), only generations of a GDG (Generation Data Group) are to be selected. Disabled by default.
HEADINGS	field column heading lines are to be included in the report. Enabled by default.
HEXBYTES	kilobyte and megabyte fields are to be based on division by 1024 (hex) rather than 1000 (decimal). Disabled by default.
IAM	FDRREPORT will attempt to identify IAM files. IAM is a separately priced product from Innovation. Disabled by default.

54.30 CONTINUED . . .

INDEX	Extract an index level from the name of every data set for sorting purposes and other uses. The index level to be extracted (1, 2, etc.) is controlled by the INDEXNUM operand of the PRINT statement. The extracted index is available as the INDEX field name. Enabled by default.
INFOMSG	failure to process or select any data sets from a specified volume results in an informative message. Enabled by default.
MAXAVAILABLE	FDREPORT will gather all information about selected data sets from all available sources. If disabled, only the sources required to provide the field names referenced in FDREPORT statements will be read. For example, the VVDS will not be read for cluster information unless you request a field which comes from the VVDS. MAXAVAILABLE is useful if you are creating an extract file for later processing. Disabled by default.
MODELSCB	use the actual name of the ABR model DSCB in any report generated from a VTOC, instead of the "***ABR" name. Disabled by default.
MODOUTPUT	when RPTYPE=DATA or COMPDATA is specified, open the output extract file (SYSUT2) with DISP=MOD (extend) in order to add data to the file. If disabled and DISP=MOD is not specified on the SYSUT2 DD statement, every PRINT with RPTYPE=DATA or COMPDATA will overwrite the extract file so that only the latest output is available. Disabled by default.
ONLINE	all online DASD volumes are to be processed. Disabled by default. If disabled, volumes specified by DISKxxxx DD statements, VOL, VOLG, and STORGRP operands, and volumes selected from the catalog, will be processed.
PAGENUMBER	report headings are to contain a page number. Enabled by default.
PCHDIAGNOS	If DIAGNOSEVVDS is also enabled, this causes FDREPORT to write the IDCAMS input necessary to correct VVDS errors to the PCHDIAG DD statement. Disabled by default.
RESERVE	VTOCs are protected from change by a RESERVE while being read. Disabled by default.
RESETSELECT	selection criteria will be reset after executing each PRINT statement. Disabled by default.
RPTDUPDSNCHK	during report generation, bypasses duplicate data set names, generating only one report line for each name. It is effective only if the data is sorted by data set name. Disabled by default.
SELTERR	set an error code if no data sets are selected (for DATATYPE=VOLDATA, if no volumes are selected). Enabled by default unless overridden in the FDR Global Option Table.

54.30 CONTINUED . . .

TAPERREAD	when DATATYPE=TVTOC is specified (to extract information from a FDR backup file), causes the entire backup to be read, collecting additional data about IAM files and directories of PDS data sets. If disabled, only the control records at the beginning of the backup are read, which generates all VTOC and VVDS related fields. Disabled by default.
TIMESTAMP	report heading is to contain a time. Disabled by default.
VVRDISPLAY	The VVDS records (VVRs for all VSAM clusters, and NVRs for non-VSAM SMS data sets) are displayed for all selected data sets. Disabled by default.

54.31 FDREPORT FIELD NAMES

FIELD NAME TABLES The following tables list the FIELDS which can be used in REPORT, SORT, (X)SELECT/ (X)EXCLUDE, and SUMMARY statements as well as PUNCH masks and TITLE statements. All of the fields listed are valid on a REPORT statement. The tables show the following information for each FIELD:

NAME – The FIELD name. On a REPORT, SORT, or SUMMARY statement, specify the FIELD NAME in the FIELD= operand. In a PUNCH mask or TITLE statement, specify it where you want its value substituted, surrounded by the field prefix/suffix characters (e.g., <DSN>). On a SELECT/ EXCLUDE statement, specify the FIELD NAME followed by an equal (=) sign and a value appropriate to the NAME (e.g., %CI=10). On a XSELECT/XEXCLUDE statement, the FIELD NAME may be followed by any of the operations supported (e.g., %CI>10)

DESCRIPTION – A brief description of the field. For fields which represent attributes or flag bytes, the strings which are used in (X)SELECT/(X)EXCLUDE statements are shown, e.g., one of the values of CATALOG is YES so you may specify CATALOG=YES. For some such fields, the value is abbreviated in the report so the abbreviated printed value is shown in parentheses, e.g., under ARCFLAG2 it says RECALL (A), so you specify ARCFLAG2=RECALL, but in a report ARCFLAG2 will include A if the recall flag is set.

LEN – The number of print positions the field will occupy on the report.

ATTR – How the field will be displayed. There are three (3) types of fields: Character (CHAR), Numeric (NUM) and Hexadecimal (HEX). MIX indicates a combination of these types.

SOURCE – The source of the information displayed.

VTOC	DSCB of the data set.
CAT	system or ABR catalog.
VVDS	VVR or NVR in SYST.VVDS.Vvolser.
ARCH	ABR archive control file.
MODL	ABR MODEL DSCB.
IAM	IAM control blocks.
TRKC	TRKCALC macro.
DIRB	PDS directory.
FAMS	SMS File Attribute Management Support.
VTIX	VTOC Index.
SMS	SMS
GEN	generated from one or more of the above sources.

Although it is not indicated under SOURCE, some FDREPORT fields are taken from the HSM MCDS or BCDS when those file are used as input. [See Section 54.32](#) for details.

PUNCH/SORT – These columns indicate whether the field is eligible for the specified statement. If an 'X' appears under the heading, it means the field can be specified on that statement. Fields valid for SORT can also be specified in a TITLE statement.

SELECT – A value in this column indicates that the field is eligible for the SELECT/EXCLUDE statements. If an 'E' appears in the column, the data set will be selected if the value in the file is EQUAL TO the value on the SELECT statement. If a 'G' appears in the column, the data set will be selected if the value is GREATER THAN or EQUAL TO the value on the SELECT statement.

XSELECT – A value in this column indicates that the field is eligible for the XSELECT/XEXCLUDE statements. If an 'E' appears in the column, only the equal (= or .EQ.) and not-equal (≠ or .NE.) comparisons may be used. If an 'A' appears, all comparisons supported by those statements may be used.

SUMMARY – A value in this column indicates that the field is eligible for the SUMMARY statement. If an 'C' (count) appears, then the number of unique occurrences of the value are simply counted. If a 'V' (value) occurs, each unique value of the field is displayed with the number of occurrences of that value. If a 'S' (sum) is displayed, then a total of all values of the field is displayed.

54.31 CONTINUED . . .

TABLE 1 lists the FIELDS available from the data set's DSCB in the VTOC. If a VSAM cluster is being reported, some of the information may be found in the VVDS or catalog.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
DSN	Data Set Name/VSAM component name	44	CHAR	VTOC	E	A	X	X	C
SPLDSN	Data Set Name on two lines	27	CHAR	VTOC					
SDSN	Data Set Name (First 20 bytes only)	20	CHAR	VTOC					
NAME	Data Set Name or VSAM Cluster Name	44	CHAR	VTOC/VVDS			X	X	
SPLNAME	Data Set Name or VSAM Cluster Name on two lines	27	CHAR	VTOC/VVDS					
SNAME	Data Set Name or VSAM Cluster Name (First 20 bytes only)	20	CHAR	VTOC					
DEFAULTS	Includes: SPLDSN,VOL,DSORG,RECFM, BLKSIZE,LRECL,SIZE,SIZEFREE,%FREE	72	MIX						
DSSN	Volume serial – stored in DSCB***	6	CHAR	VTOC	E	A		X	
VOLSQ	Data Set Volume Sequence Number	3	NUM	VTOC	E	A		X	V
DATES	Includes: CRDATE,EXPDATE,LRDATE	*	NUM						
CRDATE	Creation Date	*	NUM	VTOC/CAT	G	A	X	X	V
EXPDATE	Expiration Date	*	NUM	VTOC/CAT	G	A	X	X	V
LRDATE	Last Reference Date	*	NUM	VTOC/VVDS	G	A	X	X	V
NOEPV	Number of Extents for Data Set	3	NUM	VTOC	G	A	X	X	S
NOBDB	Number of Bytes in Last Directory Block	3	NUM	VTOC	G	A	X	X	
SYSCODE	System Code Field in DSCB***	13	CHAR	VTOC	E	A		X	
DCBINFO	Includes: DSORG,RECFM,BLKSIZE,LRECL	21	MIX						
DSORG	Data Set Organization: DA-BDAM PS- SEQUENTIAL U- UNMOVABLE AM- Non-ICF VSAM PO- PARTITIONED UN- UNDEFINED EF- ICF VSAM POE- PDS PSE-PS Extended Format (Striped)IAM-Innovation Access Method HFS-Hierarchical File System (Open Edition)	3	CHAR	VTOC	E	E	X	X	V
RECFM	Record Format	5	CHAR	VTOC	E	E	X	X	V
OPTCD	Option Code Byte	2	HEX	VTOC	E	E	X	X	V
BLKSIZE	Data Set Block Size	5	NUM	VTOC/VVDS	E	A	X	X	V
LRECL	Data Set Logical Record Size	5	NUM	VTOC/VVDS	E	A	X	X	V
KEYLEN	Data Set Key Length	3	NUM	VTOC/VVDS	E	A	X	X	V
RKP	Data Set Relative Key Position	4	NUM	VTOC/VVDS	E	A	X	X	V
DSIND	Data Set Indicators - printed in hex; for selection use: LASTV - last volume (X'80') UPDAT - data set updates (X'02') RACF -discrete profile (X'40') PASSW - write password (X'14') PASSA -read/write password (X'10')	2	HEX	VTOC	E	E		X	V
SECALLOC	Secondary Allocation Quantity	8	NUM	VTOC/VVDS	E	A	X	X	V
SECAFLAG	Secondary Allocation Flags: BLK - allocated in blocks FIV - 5 largest extents (ALX) CYL - allocated in cylinders RND - rounded to cylinders TRK - allocated in tracks (ROUND) CON - contiguous (CONTIG) MAX - largest extent (MXIG)	5	CHAR	VTOC/VVDS	E	A	X	X	V
SECXFLAG	Secondary Allocation extension flag	3	CHAR	VTOC				X	V
SECXVALU	Secondary Allocation extension value	11	NUM	VTOC				X	V
LSTAR	Data Set Last Block Pointer; specify as LSTAR=X'ttttr'	6	HEX	VTOC	E	A		X	
TRKBAL	Data Set Track Balance; remaining bytes on last track used	5	NUM	VTOC	G	A		X	
EXTENTS	Data Set Extents; each extent will be stacked; up to 123 extents may be printed	15**	CHAR	VTOC		A			
EXTENTSX	Data Set Extents (with extent type); each extent will be stacked, up to 123 extents may be printed	17**	CHAR	VTOC					
CTFLD	Count Field of data set DSCB (CCHHR)	10	HEX	VTOC		A	X	X	
BPTR	Format 1 DSCB Pointer to Format 2/3 DSCB	10	HEX	VTOC				X	
LMJOB	Last Modifying JOB (for ASM2 users)***	8	CHAR	VTOC		A		X	

* the format and width of date fields varies depending on the value of the DATEFORMAT= operand.

** these fields will be two columns larger if the FATDISK option is enabled.

*** this field supports selection via a mask; see the VOL= operand in Section 54.13

54.31 CONTINUED . . .

TABLE 2 lists the FIELDS available from the VVDS or ICF catalog for ICF VSAM clusters. This is in addition to the FIELDS available from the VTOC (TABLE 1). For non-VSAM data sets, these fields will contain blanks. Those fields that show a source of IAM are also valid for IAM files; those that show IAM+ are valid only for enhanced IAM files.					SELECT	XSELECT	PUNCH	SORT	SUMMARY
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
CLUSTER	ICF VSAM Cluster Name	44	CHAR	VVDS	E	A	X	X	V
SPLCLS	ICF VSAM Cluster Name on two lines	27	CHAR	VVDS					
SCLUSTER	Cluster name (First 20 bytes only)	20	CHAR	VVDS					
VSAMID	Includes: SPLCLS,SPLDSN,VOLSER	62	CHAR						
VSAMINFO	Includes: DSORG,RECFM,MAXLRECL,LRECL,CISIZE,TRKSCA,CICA,%CI,%CA,COMPATTR	68	MIX						
VSAMUSE	Includes: RECORDS,RETRIEVE,UPDATES,INSERTS,DELETES,HIALORBA,HIUSERBA,CISPLIT,CASPLIT	85	NUM						
TIMES	Includes: CRTIME,LRTIME	17	NUM						
CRTIME	Time of creation (hh.mm.ss)	8	NUM	VVDS/IAM				X	
LRTIME	Time of last reference (hh.mm.ss)	8	NUM	VVDS/IAM				X	
CATNAME	Name of the Catalog	44	CHAR	VVDS/IAM	E	A		X	V
%CA	Free Percent in Control Area (CA)	3	NUM	VVDS/IAM	G	A	X	X	V
%CI	Free Percent in Control Interval (CI)	3	NUM	VVDS/IAM	G	A	X	X	V
CICA	Number of Control Intervals per CA	3	NUM	VVDS	G	A	X	X	V
CISIZE	Control Interval (CI) Size	5	NUM	VVDS/IAM	E	A	X	X	V
BUFSIZE	VSAM Buffer Size	6	NUM	VVDS	E	A	X	X	V
MAXLRECL	VSAM Maximum Record Length	6	NUM	VVDS/IAM	E	A	X	X	V
TRKSCA	VSAM Number Tracks per Control Area	3	NUM	VVDS	G	A	X	X	
CLUSATTR	VSAM Cluster Attributes ATL – Tape Volume Catalog ATTREXT – Extended Attribute (EA) EXTEND – Extended Format (EF) FORMAT – Formatted Pagespace PAGESPACE – Pagespace SWAPSPACE – Swapspace TIMESTAMP – Timestamps Exist VERIFYREQ – Verify is Required	24	CHAR	VVDS		E		X	
COMPTYPE	VSAM Component type: DATA – Base cluster data INDEX – Base cluster index AIXDATA (AIXDA) – Alternate Index data AIXINDEX (AIXIN) – Alternate Index index	5	CHAR	VVDS		E		X	
COMPATTR	VSAM Component Attributes IXD – Indexed RUS – Reusable KRNG – Key Range NUMD – Relative RCVY – Recovery SPED – Speed IMBD – Imbedded NIXD – non-Indexed ERAS – Erase ORD – Ordered WCK – Write Check SPND – Spanned REPL – Replicate	24	CHAR	VVDS			X		
AMDATTR	VSAM Cluster Attributes from AMDSB: ESDS (ES) IMBED (IM) KEYRANGE (KE) KSDS (KS) ORDERED (OR) RRDS (RR) REPLICATE (RE) SPANNED (SP) WRITECHECK (WR)	24	CHAR	VVDS/IAM		E		X	
AMDATTR3	VSAM Cluster Attributes from AMDSB: LINEAR (LI) LOADED (LO) NONUNIQUE (NO) SHRBCS (SH) – shared catalog VARIABLE (VA) - Variable RRDS	24	CHAR	VVDS/IAM		E		X	
AIXATTR	Attribute of Alternate Index: AIX – this is an AIX UPGRADE - AIX with the UPGRADE attribute	7	CHAR	VVDS		E		X	V
SHROPT	VSAM Share Options	3	CHAR	VVDS/IAM			X		V
PRIALLOC	VSAM Primary Allocation Quantity	8	NUM	VVDS/IAM	E	A	X	X	V
HIALORBA	VSAM High Allocated Relative Byte Address	10	NUM	VVDS/IAM+	G	A	X	X	
HIKEYRBA	VSAM High Key Relative Byte Address	10	NUM	VVDS/IAM+	G	A	X	X	
HIUSERBA	VSAM High Used Relative Byte Address	10	NUM	VVDS/IAM+	G	A	X	X	
EXCPEXIT	VSAM Exception Exit *	8	CHAR	VVDS	E	A			V
OWNER	VSAM Owner Identification	8	CHAR	CAT	E	A	X	X	V
TIMESTMP	VSAM Time Stamp	16	HEX	VVDS/IAM				X	
CASPLIT	Number of Control Area (CA) Splits	5	NUM	VVDS	G	A	X	X	S
CISPLIT	Number of Control Interval (CI) Splits	5	NUM	VVDS	G	A	X	X	S
INDEXLEV	VSAM Number of Index Levels	5	NUM	VVDS		A	X	X	V
EXCPS	VSAM Number of EXCPs	10	NUM	VVDS	G	A	X	X	S
DELETES	VSAM Number of Records Deleted	10	NUM	VVDS/IAM	G	A	X	X	S
INSERTS	VSAM Number of Records Inserted	10	NUM	VVDS/IAM	G	A	X	X	S
RECORDS	VSAM Number of Records in Cluster	10	NUM	VVDS/IAM	G	A	X	X	S
RETRIEVE	VSAM Number of Records Retrieved	10	NUM	VVDS	G	A	X	X	S
UPDATES	VSAM Number of Records Updated	10	NUM	VVDS/IAM	G	A	X	X	S
VSFREBYT	VSAM Number of Bytes Free	10	NUM	VVDS	G	A		X	S
CATVRBA	Relative Byte Address of VVR from catalog	8	HEX	CAT	G	A	X	X	V
VVRVRBA	Actual Relative Byte Address of VVR in VVDS	8	HEX	VVDS	G	A		X	V

* this field supports selection via a mask; see the VOL= operand in Section 54.13

54.31 CONTINUED . . .

TABLE 3 lists the additional FIELDS available for a data set on an SMS-managed volume. For non-SMS data sets, these fields will be blank.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
SMSCLASS	Includes: STORCLAS,DATACLAS,MGMTCLAS	26	CHAR						
STORCLAS	SMS Storage Class*	8	CHAR	VVDS		E	X	X	V
MGMTCLAS	SMS Management Class*	8	CHAR	VVDS		E	X	X	V
DATACLAS	SMS Data Class*	8	CHAR	VVDS		E	X	X	V
STORGRP	SMS Storage Group on which the data set resides*	8	CHAR	GEN		E	X	X	V
SMSFLAGS	SMS flag byte in DSCB: MANAGED (S) - SMS-managed NOBCS (N) - not cataloged REBLOCK (R) - reblockable DADSMCRT (D) - DADSM assigned blksize PDSE (I) - PDS Extended PDSEX (X) - HFS (Hierarchical File Sys) STRIPE (Z) - Extended Format (EF) ATTREXTN (E) - Extended Attribute (EA)	4	CHAR	VTOC		E		X	V
LASTBKUP	SMS last backup Date and Time (HSM)	13	CHAR	VVDS		A		X	
NVSAMFLG	SMS NONVSAM Flags: ATTREXT- Extended Attribute (EA) EXTEND – Extended Format (EF)	7	CHAR	VVDS		E		X	V
NVSAMATR	SMS NONVSAM Attributes: ACTGDG - active GDG generation DEFGDG - deferred GDG generation ROLGDG - rolled-out GDG generation PDSE - PDS Extended POSIX - HFS (Hierarchical File Sys)	7	CHAR	VVDS		E		X	V
STRIPECT	Stripe Count, for Extended Format (EF)	3	CHAR	VVDS				X	V
RECOVDTA	SMS Recovery Data	8	CHAR	VVDS		E		X	V
RESOROWN	SMS Resource Ownership	16	CHAR	VVDS		E		X	V
SUBCFLAG	SMS Subcell Flag Byte	17	CHAR	VVDS		E		X	V
SUBCVERS	SMS Subcell Version Number	3	NUM	VVDS		A		X	V
COMPDATA	Compress Data Set Size	10	NUM	VVDS		A	X	X	S
USERDATA	Original Size of Compressed Data Set	10	NUM	VVDS		A	X	X	S

* this field supports selection via a mask; see the VOL= operand in Section 54.13

54.31 CONTINUED . . .

Table 4 lists data set backup information from the ABR system. If DATATYPE=ARCHIVE was specified, this information is from the archive control file, otherwise it is from the ABR backup subsystem. If COPY= or OLDBACKUP= was specified, there may be more than one line of this information available; REPORT will list them on multiple lines, SELECT/XSELECT will test against all copies of the fields.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
ABRGEN	Current ABR Generation Number	4	NUM	MODL	E	A	X	X	V
ARBCYCLE	Current ABR Cycle Number	3	NUM	MODL	E	A	X	X	V
ABRIND	ABR Indicators: A - Always backup/never archive B - Current ABR backup exists N - Normal backup/never archive R - Archive requested X - Exclude from ABR processing For reports only, not for selection: C - No current ABR backup exists M - Multi-volume data set T - old backup recording is enabled U - Updated since last ABR backup	5	CHAR	VTOT	E	E		X	V
ADATE	Archive Date	*	NUM	ARCH	G	A	X	X	V
ATIME	Archive time (hhmmss, Application Bkup)	8	NUM	ARCH	G	A	X	X	V
ADAYS	# days since Archive	5	NUM	ARCH	G	A	X	X	V
ARCDSN	Archive Control File name	44	CHAR	GEN			X		V
ARCFLAGS	Includes: ARCFLAG1, ARCFLAG2								
ARCFLAG1	ARCHIVE control file flag 1 - FIVEVOL (F) - backup over 5 volumes MULTIVOL (M) - multi-vol data set RESTORED (R) - restored from ARCHIVE	5	CHAR	ARCH		E		X	V
ARCFLAG2	ARCHIVE control file flag 2 - CLUSTER (C) - entry for ICF cluster DELETE (D) - entry flagged for deletion NOTCAT (N) - backup not cataloged RECALL (A) - archived for auto-recall	5	CHAR	ARCH		E		X	V
ARCTTR	TTR used for auto recall	6	HEX	ARCH		A	X	X	V
ACTTTR	Actual TTR of record in ARCHIVE ctl file	6	HEX	ARCH		A	X	X	V
BKINFO	Includes: BKDATE,BKSUFFIX,BKFILENO,BKVOL	*	CHAR						
TVTOCDSN	Input Backup dsn, for RPTYPE=TVTOC	44	CHAR	GEN			X	X	V
OLDBKUP	ABR Old Backup # of ABR backup displayed	2	NUM	GEN				X	V
BKCYCLE	Cycle # of ABR backup displayed	3	NUM	CAT		A	X	X	V
BKGEN	Generation # of ABR backup displayed	4	NUM	CAT		A	X	X	V
BKDATE	Backup Date	*	NUM	CAT	G	A	X	X	V
BKDAY	# days since Backup	5	NUM	CAT	G	A	X	X	V
BKEXDATE	Backup or Archive expiration date	*	NUM	ARCH/CAT		A	X	X	V
BKEXDAY	# days until expiration of Backup or Archive**	5	NUM	ARCH/CAT		A	X	X	V
BKFILENO	File Seq Number of Backup or Archive	4	NUM	ARCH/CAT	E	A	X	X	V
BKSUFFIX	Backup or Archive Data Set Name Suffix. This is the last index level in the name***	8	CHAR	ARCH/CAT	E	A	X	X	V
BKTAPCNT	Number of Backup or Archive Volumes	2	NUM	ARCH/CAT	E	A	X		V
BKVOL	Volume Serials for the Backup/Archive (5)***	34	CHAR	ARCH/CAT	E	A	X	X	V
BKXVOL	Volume Serials for the Backup/Archive (20)	139	CHAR	ARCH/CAT	E	A			V
BKDEVCLS	Device class of backup (TAPE or DISK)	4	CHAR	ARCH/CAT			X	X	V
BKDEVTP	Device type of backup	7	CHAR	GEN		E	X	X	V

* the format and width of date fields varies depending on the value of the DATEFORMAT= operand.

** if the expiration date is 99.000 or 99.365 or above, this is set to 65535.

*** this field supports selection via a mask; see the VOL= operand in Section 54.13

54.31 CONTINUED . . .

TABLE 5 lists the FIELDS which can be generated by ABR from information contained in the VTOC, Catalog or the VVDS.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
SIZEINFO	Includes: SIZE,SIZEFREE,%FREE	15*	NUM						
SIZE	Size of the Data Set in Tracks	5*	NUM	GEN	G	A	X	X	S
BYTES	Size of the Data Set in bytes/KB/MB	**	NUM	GEN		A	X	X	S
MAXSIZE	Maximum size of Data Set in Tracks if all secondary allocations taken	5*	NUM	GEN	G	A	X	X	S
SIZEFREE	Number of Tracks Unused for Data Set – PS, PO and VSAM actual free tracks – all others zero (0)	5*	NUM	GEN	G	A	X	X	S
BYTESFRE	Bytes unused in the Data Set	**	NUM	GEN		A	X	X	S
SIZEUSED	Number of Tracks Used for Data Set – PS, PO and VSAM actual used tracks – all others total allocated space	5*	NUM	GEN	G	A	X	X	S
BYTESUSE	Bytes used in the Data Set	**	NUM	GEN		A	X	X	S
BLKSTRK	Number of Blocks per Track	5	NUM	VVDS/TRKC	G	A	X	X	V
BYTESTRK	Bytes per Track (BLKSIZE times BLKSTRK)	5	NUM	GEN		A	X	X	V
PRIBYTES	Bytes in primary allocation (ICF VSAM only)	**	NUM	GEN		A	X	X	S
SECBYTES	Bytes/KB/MB size of secondary allocation	**	NUM	GEN		A	X	X	S
TRACKCAP	Max Track Capacity of Device in Bytes	5	NUM	VTOC			X	X	V
CAPBYTES	Bytes size of allocated space if used at track capacity	**	NUM	GEN		A	X	X	S
%FREE	Percentage of Free Space in Data Set	3	NUM	GEN	G	A	X	X	V
%USED	Percentage of Used Space in Data Set	3	NUM	GEN	G	A	X	X	V
%CAPUSED	Percentage of Capacity Utilization (based on allocation and BLKSIZE)	3	NUM	GEN		A	X	X	V
CATALOG	Indicates if Data Set is Cataloged This can be expensive to collect if a large number of data sets are to be reported. YES – cataloged to this volume NO – not cataloged at all ERR – cataloged to another volume ONL – only cataloged, not in VTOC UNK – error reading catalog CAN – cataloged to candidate volume DRF – disk read failure	3	CHAR	CAT	E	E		X	V
CATTTR	DSCBTTR from Catalog (used by ABR to indicate auto recall)	6	HEX	CAT		A	X	X	X
CATVOL	Volume to which data set is cataloged	6	CHAR	CAT			X	X	X
CATVOLCT	Number of volumes to which the data set is cataloged	3	NUM	CAT		A	X	X	V
CRDAYS	Number of days since creation	5	NUM	GEN	G	A	X	X	V
EXPDAYS	Number of days until expiration****	5	NUM	GEN	G	A	X	X	V
LRDAYS	Number of days since last reference	5	NUM	GEN	G	A	X	X	V
NOEXTENT	Number of Extents Used for Data Set	3	NUM	GEN	G	A	X	X	S
PROTECT	Reports on Type of Protection indicates if RACF or PASSWORD protected	4	CHAR	VTOC				X	V
RUNDATE	Date of FDREPORT execution	**	NUM	GEN		A	X	X	V
RUNTIME	Time of FDREPORT execution (hh.mm.ss)	8	NUM	GEN		A	X	X	
VOL	Volume Serial Data Set resides on*****	6	CHAR	VTOC	E	A	X	X	V
FILESEQ	File Sequence Number (tape only)	4	NUM	CAT		A	X	X	
UNIT	Device address the data set is on	3/4	CHAR	GEN		A	X	X	V
DEVTYPE	Type of Device Data Set resides on	7	CHAR	GEN	E	E	X	X	V
DEVCLASS	Class of Device Data Set resides on, "TAPE" or "DISK"	4	CHAR	GEN			X		V
UCBID	4-byte UCB device type	8	HEX	CAT/UCB			X		V
INDEXNUM	# of Index Level in the data set or cluster name that is extracted into field INDEX	3	NUM	USER			X	X	V
INDEX	Extracted level from data set or Cluster name	8	CHAR	DSN/CLUS			X	X	V

54.31 CONTINUED . . .

					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
TABLE 5 lists the FIELDS which can be generated by ABR from information contained in the VTOC, Catalog or the VVDS.									
NTMIGRAT	HSM only: number of times data set has been migrated (MCDS) or backed up (BCDS)	8	CHAR		E	E		X	
SOURCE	Source of this data record (ARCHIVE, CATALOG, VTOC, VVDS, MCDS, BCDS)	8	CHAR		E	E		X	

- * the width of the size (in tracks) fields will be 6 if the FATDISK option is enabled.
- ** the format and width of byte fields varies depending on the value of the BYTEFORMAT= operand.
- *** the format and width of date fields varies depending on the value of the DATEFORMAT= operand.
- **** if the expiration date is 99.000 or 99.365 or above, this is set to 65535.
- ***** this field supports selection via a mask; see the VOL= operand in Section 54.13

54.31 CONTINUED . . .

TABLE 6 lists the FIELDs which can be obtained from the directory of a Partitioned Data Set or an Extended Partitioned Data Set (via FAMS).					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
DIRBINFO	Includes: DIRBLOCK,DIRBFREE,%DIRFREE		NUM						
DIRBLOCK	Number of PDS Directory Blocks	5	NUM	DIRB/FAMS		A	X	X	V
DIRBFREE	Number of Free PDS Directory Blocks	5	NUM	DIRB/FAMS		A	X	X	V
DIRBUSED	Number of used PDS Directory Blocks	5	NUM	DIRB/FAMS		A	X	X	V
%DIRFREE	Percentage of PDS Directory Blocks not used	3	NUM	GEN		A	X	X	V
%DIRUSED	Percentage of PDS Directory Blocks in use	3	NUM	GEN		A	X	X	V
MEMBERS	Number of members in PDS	6	NUM	DIRB/FAMS		A	X	X	S

TABLE 7 lists the FIELDs available from IAM data set Control Block for IAM datasets. This is in addition to the FIELDs available from the VTOC (TABLE 1) and the fields that are normally reserved for VSAM.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
IAMINFO	Includes: DSORG,RECFM,MAXLRECL, LRECL, BLKSIZE, CISIZE	33	MIX						
IAMUSE	Includes: RECORDS,UPDATES,INSERTS, DELETES,OVERFLOW,OVERUSED,%OVER	69	NUM						
IAMINDIC	IAM Indicators: ENHANCED (E) - Enhanced IAM format DATACOMP (D) - Data compressed KEYCOMP (K) - Keys compressed	5	CHAR	IAM		E		X	V
%PRIMEXT	Used Percent of the Prime Extension	3	NUM	IAM		A	X	X	V
PRIMEXTN	Number of Prime Extension Blocks allocated	8	NUM	IAM		A	X	X	S
PRIMEUSE	Number of used Prime Extension Blocks	8	NUM	IAM		A	X	X	S
%OVER	Used % of the Independent Overflow records	3	NUM	IAM		A	X	X	V
OVERFLOW	# of Independent Overflow records allocated	10	NUM	IAM		A	X	X	S
OVERUSED	# of used Independent Overflow records	10	NUM	IAM		A	X	X	S

54.31 CONTINUED . . .

TABLE 8 lists the FIELDS which can be obtained from the volume-level record. Data is gathered and summarized from the LSPACE SVC, VTOC, VTOC index, and VVDS.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
VLDEFAULT	Volume report defaults, includes: VOL, DEVTYPE, UNIT, VLINDSTA, VLUSEATR, VLTRKVOL, VLUSERS, VL%UTRKS, VLFRETRK, VLLRGCYL, VLLRGTRK, VLFREEXT and VLFRAGIN		CHAR		E	A	X	X	V
VLVOLSER	Volume Serial Data Set resides on ***	6	CHAR	VTOC	E	A	X	X	V
VLUNIT	Device address the volume is on	3/4	CHAR	GEN		E	X	X	V
VLDEVTYPE	Type of Device Data Set resides on	7	CHAR	GEN	E	E	X	X	V
VL%FTRKS	Percentage of disk volume free	3	NUM	VTOC		A	X	X	V
VL%UTRKS	Percentage of disk volume in use	3	NUM	VTOC		A	X	X	V
VL%FDSCB	Percentage of DSCBs free within the VTOC	3	NUM	VTOC		A	X	X	V
VL%UDSCB	Percentage of DSCBs used within the VTOC	3	NUM	VTOC		A	X	X	V
VL%FINDX	Percentage of VTOC Index Records free	3	NUM	VTIX		A	X	X	V
VL%UINDX	Percentage of VTOC Index Records used	3	NUM	VTIX		A	X	X	V
VL%FVVDS	Percentage of VVDS Control Intervals free	3	NUM	VVDS		A	X	X	V
VL%UVVDS	Percentage of VVDS Control Intervals used	3	NUM	VVDS		A	X	X	V
VLALOCYL	Number of Allocated Cylinders	5	NUM	VTOC		A	X	X	S
VLALOTRK	Number of Allocated Tracks	5*	NUM	VTOC		A	X	X	S
VLALOBYT	Number of Allocated Bytes	**	NUM	VTOC		A	X	X	S
VLALTRKS	Number of Alternate Tracks on the volume	5	NUM	DEVT		A	X	X	S
VLAVGTRK	Average File Size in Tracks	5*	NUM	VTOC		A	X	X	V
VLAVGBYT	Average File Size in Bytes	**	NUM	VTOC		A	X	X	V
VLBYTRK	Track Capacity in bytes	5	NUM	DEVT			X	X	V
VLCYLVOL	Number of Cylinders per Volume	5	NUM	DEVT			X	X	S
VLTRKCYL	Number of Tracks per Cylinder	3	NUM	DEVT			X	X	V
VLBYTCYL	Number of Bytes per Cylinder	**	NUM	DEVT			X	X	V
VLTRKVOL	Number of Tracks per Volume	5*	NUM	DEVT			X	X	S
VLBYTVOL	Number of Bytes per Volume	**	NUM	DEVT			X	X	S
VLDIRBTR	Number of PDS Directory Blocks per Track	3	NUM	DEVT			X	X	V
VLDSCBTR	Number of VTOC DSCBs per Track	3	NUM	DEVT			X	X	V
VLDSCB	Total Number of DSCBs	5*	NUM	VTOC		A	X	X	S
VLDSCBU	Number of Used DSCBs	5*	NUM	VTOC		A	X	X	S
VLDSCB0	Number of FORMAT-0 DSCBs	5*	NUM	VTOC		A	X	X	S
VLDSCB1	Number of FORMAT-1 DSCBs	5*	NUM	VTOC		A	X	X	S
VLDSCB2	Number of FORMAT-2 DSCBs	5*	NUM	VTOC			X	X	S
VLDSCB3	Number of FORMAT-3 DSCBs	5*	NUM	VTOC		A	X	X	S
VLDSCB4	Number of FORMAT-4 DSCBs	5	NUM	VTOC			X	X	S
VLDSCB5	Number of FORMAT-5 DSCBs	5	NUM	VTOC			X	X	S
VLDSCB6	Number of FORMAT-6 DSCBs	5	NUM	VTOC			X	X	S
VLDSCB7	Number of FORMAT-7 DSCBs	5	NUM	VTOC			X	X	S
VLDSOAM	Number of non-ICF VSAM Data Spaces	5*	NUM	VTOC		A	X	X	S
VLDSODA	Number of Direct Access data sets	5*	NUM	VTOC		A	X	X	S
VLDSOEF	Number of ICF/VSAM Components	5*	NUM	VTOC		A	X	X	S
VLDSOIS	Number of Indexed Sequential data sets	5*	NUM	VTOC		A	X	X	S
VLDSOPO	Number of Partitioned data sets	5*	NUM	VTOC		A	X	X	S
VLDSOPOE	Number of PDSE data sets (SMS volumes)	5*	NUM	VTOC		A	X	X	S
VLDSOPS	Number of Physical Sequential data sets	5*	NUM	VTOC		A	X	X	S
VLDSOUN	Number of data sets with an Unknown data set Organization	5*	NUM	VTOC		A	X	X	S
VLFRAGIN	IBM Fragmentation Index	4	NUM	LSPC		A	X	X	V
VLFRECYL	Number of Unused Cylinders	5	NUM	LSPC		A	X	X	S
VLFREBYT	Number of Unused Bytes	**	NUM	VTOC			X	X	S

54.31 CONTINUED . . .

TABLE 8 lists the FIELDS which can be obtained from the volume-level record. Data is gathered and summarized from the LSPACE SVC, VTOC, VTOC index, and VVDS.					S E L E C T	X S E L E C T	P U N C H	S O R T	S U M M A R Y
NAME	DESCRIPTION	LEN	ATTR	SOURCE					
VLFREEEXT	Number of Unused Extents	5*	NUM	LSPC		A	X	X	S
VLFRETRK	Number of Unused Tracks	5*	NUM	VTOC		A	X	X	S
VLFREVC I	Number of Unused VVDS Control Intervals	5	NUM	VVDS		A	X	X	S
VLUSEVC I	Number of Used VVDS Control Intervals	5	NUM	VVDS			X	X	S
VLTOTVC I	Total Number of VVDS Control Intervals	5	NUM	VVDS			X	X	S
VLFREVIR	Number of Unused VTOC Index Records	5	NUM	LSPC		A	X	X	S
VLUSEVIR	Number of Used VTOC Index Records	5	NUM	GEN			X	X	S
VLTOTVIR	Total Number of VTOC Index Records	5	NUM	GEN			X	X	S
VLALOSTA	Volume Allocation Status: SYSRES - system residence volume ALLOC - allocated UNLOAD - unload pending ONLINE - online, not allocated or sysres	6	CHAR	UCB				X	X
VLINDSTA	Status of Indexed VTOC ACTIVE – Indexed VTOC Active NONE – No Indexed VTOC YES – Indexed VTOC Not Active	6	CHAR	LSPC		E		X	V
VLMOUSTA	Volume Mount Status RESERVED – Volume Dismountable RESIDENT – Volume Not Dismountable	8	CHAR	UCB				X	V
VLUSEATR	Volume Use Attribute PRIVATE – Allocated if specific PUBLIC – Temporary non-specific STORAGE – Non-temporary non-specific	7	CHAR	UCB		E		X	V
VLMSSTAT	Volume SMS State INITIAL – Volume Initialized for SMS MANAGED – Volume SMS Managed NONE – No SMS Processing	7	CHAR	VTOC		E		X	V
VLMSVST	Volume SMS Status ENABLED NONE (non-SMS) QUIESCED ALL QUIESCED NEW DISABLED ALL DISABLED NEW	12	CHAR	SMS				X	V
VLMSGST	Volume SMS Storage Group Status (same values as VLMSVST)	12	CHAR	SMS				X	V
VLLRGCYL	Number of Cylinders in Largest Free Extent	5	NUM	LSPC		A	X	X	S
VLLRGTRK	Number of Tracks in Largest Free Extent	5*	NUM	LSPC		A	X	X	S
VLLRGBYT	Number of Bytes in Largest Free Extent	* **	NUM	LSPC		A	X	X	S
VLUSERS	Number of Current Allocations to Volume	5	NUM	UCB		A		X	S
VLVTIXTR	Number of Tracks Allocated to the VTOC Index	5	NUM	VTOC		A		X	S
VLVTOCTR	Number of Tracks Allocated to VTOC	5	NUM	VTOC		A		X	S
VLVVDSTR	Number of Tracks Allocated to VVDS	5	NUM	VTOC		A		X	S
VLVVDSXT	Number of Extents Allocated to VVDS	3	NUM	VTOC		A		X	S

* the width of these fields will be 6 if the FATDISK option is enabled.

** the format and width of byte fields varies depending on the value of the BYTEFORMAT= operand and the FATDISK option.

*** this field supports selection via a mask; see the VOL= operand in Section 54.13

54.32 FDREPORT HSM REPORTING

**HSM
REPORTING**

FDREPORT customers who are also using IBM's DFSMSHsm DASD management system (or its older version DFHSM), can use FDREPORT can read the data base files created by HSM and report on their contents. The HSM data sets supported by FDREPORT are:

MCDS - information on migrated data sets

BCDS - information on backups of current data sets

This allows you to use all the power and flexibility of FDREPORT to generate reports about the data that HSM is managing.

With the exception of NTMIGRAT, there are no new field names defined when reporting from HSM data. Instead, data fields from the HSM records and calculated values are mapped into standard FDREPORT field names.

For DATATYPE=MCDS, the MCDS data set record is described by the IBM MCD macro and the mapping is:

MCDS field name	FDREPORT field name	Field Description
MCK	DSN	Data set name
MCDFRVSN	VOL	Volume Serial Number
MCDUCBTY	UCBID	UCB 4-byte device type(hex)
	DEVTYPE	Disk device type (e.g., 3390)
MCDFLGS	DSSN	MGRATx (x=migration level)
MCDDLCL	CRDATE	Creation Date
	CRDAYS	Days since Creation (calculated)
MCDEXPDT	EXPDATE	Expiration Date
	EXPDAYS	Days since Expiration (calculated)
MCDDLRL	LRDATE	Last Reference Date
MCDTLRL	LRTIME	Last Referenced Time
	LRDAYS	Days since Last Reference (calculated)
MCDDMIG	ADATE	Date Archived (Migrated)
MCDTMIG	ATIME	Time Archived (Migrated)
	ADAYS	Days since Migration (calculated)
MCDDSORG	DSORG	Data Set Organization
MCDRECFM	RECFM	Record Format
MCDOPTCD	OPTCD	Option Code
MCDBLKSZ	BLKSIZE	Block Size
	BLKSTRK	Blocks/track (calculated)
	BYTESTRK	Bytes/track (calculated)
MCDKEYLN	KEYLEN	Key Length
MCDDSIND	DSIND	Data Set Indicators
MCDSIZE	SIZE	Allocated Tracks
	SIZEUSED	Used Tracks (calculated)
	SIZEFREE	Free Tracks (calculated)
MCDSIZEB	BYTES	Bytes Allocated
	BYTESUSE	Bytes Used (calculated)
	BYTESFRE	Bytes Free (calculated)
	%USED	Percentage used (calculated)
	%FREE	Percentage free (calculated)
	LSTAR	Last Used ttttrr
MCDSCNAM	STORCLAS	SMS Storage Class
MCDMCNAM	MGMTCLAS	SMS Management Class
MCDMDNAM	DATACLAS	SMS Data Class
MCDSMSFG	SMSFLAGS	SMS Data Set Flag Byte
MCDSCAL1	SECAFLAG	Secondary Allocation Flags
MCDSCAL3	SECALLOC	Secondary Allocation Quantity
MCDNMIG	NTMIGRAT	Number of times data set was migrated
	SOURCE	'MCDS'

54.32 CONTINUED . . .

For DATATYPE=BCDS, the BCDS data set record is described by the IBM MCB macro and the mapping is:

BCDS field name	FDREPORT field name	Field Description
MCK	DSN	Data set name
MCBFRVOL	VOL	Volume Serial Number
MCBDBU	ADATE	Date Backed Up
MCBTBU	ATIME	Time Backed Up
	ADAYS	Days since backup (calculated)
MCBDLRPD	LRDATE	Last Reference Date
	LRDAYS	Days since last reference (calculated)
MCBDSORG	DSORG	Data Set Organization
MCBRECFM	RECFM	Record Format
MCBBLKSZ	BLKSIZE	Block Size
MCBKEYLN	KEYLEN	Key Length
MCBDSIND	DSIND	Data Set Indicators
MCBSIZE	SIZE	Allocated Tracks
MCBSIZEB	BYTESUSE	Bytes Used
	BYTES	Bytes Allocated (calculated)
MCBNBC	NTMIGRAT	Number of times data set was backed up
	SOURCE	'BCDS'

54.40 FDREPORT VTOC EXAMPLES

This section shows examples which read the VTOCs of selected volumes directly (DATATYPE=VTOC, which is the default).

A sample of the generated report is shown after each example; because of space limitations it may be condensed. In examples where SORTing is required, SORTALLOC=YES has been specified to dynamically allocate required SORT files; in your installation you may have to provide SORT JCL.

REPORT ON PREFIX FDREPORT is to report on all data sets starting with the hi-level index "ABC" on any online volume. The default report (See Section 54.14) is to be printed. No sorting is to be done; a separate page will be generated for data sets selected from each disk volume. Messages about volumes for which no data sets were selected are suppressed. The report will be done to SYSPRINT.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='ABC DATA SETS ON VOLUME <VOL>'
              XSELECT  XDSN=ABC.*
              PRINT    ENABLE=ONLINE,DISABLE=INFOMSG
```

The generated report will look like:

ABC DATA SETS ON VOLUME ABC123

DATA SET NAME	VOLSER	ORG	RECFM	BLKSZ	LRECL	ALLOC	FREE	%FR
ABC.FINDMAC.LIST	ABC123	PS	FB	3120	80	1	0	0
ABC.JCL.CNTL	ABC123	PO	FB	6160	80	5	2	40

REPORT TSO LIST DATA SETS FDREPORT is to report on all data sets residing on online volumes starting with "TSO" whose last qualifier is "LIST" or "OUTLIST" and which have not been referenced in 2 or more days. Only the data set name, volume, and date of last use are to be printed. A standard summary is to be printed for each volume processed.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//ABRSUM      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='TSO LIST AND OUTLIST DATA SETS'
              XSELECT  XDSN=(**.*LIST,**.*OUTLIST),
                      VOLG=TSO,LRDAYS>1
              REPORT   FIELD=(DSN,VOL,LRDATE)
              PRINT    SUM=YES
```

The generated report will look like:

TSO LIST AND OUTLIST DATA SETS

DATA SET NAME	VOLSER	LRDATE
USER1.SPFTEMP1.LIST	ABC123	96.134
USER3.PRT.OUTLIST	ABC321	96.257

The summary will look like:

```
SUBTOTAL -- VOL -----ABC123
DSN----5 NOEXTENT----8 SIZE----75 SIZEFREE---12 SIZEUSED---62
FINAL TOTALS
DSN----56 NOEXTENT--117 SIZE-- 352 SIZEFREE--106 SIZEUSED--246
```

54.40 CONTINUED . . .

VSAM REPORT FDREPORT is to report on selected ICF VSAM clusters on all online volumes whose serial number starts with 'PROD'. A variety of VSAM fields will be reported (many others are available).
 ENABLE=AUTOSTACK allows FDREPORT to stack fields with like attributes in order to fit the fields within the pagewidth of 80 characters. All selected clusters will be combined into one report, sorted on cluster name and component name within cluster.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='VSAM INFO FOR SELECTED CLUSTERS'
              SELECT  DSG=(IPCS,MVS,M211),DSORG=EF,VOL=PROD*
              REPORT  FIELD=(CLUSTER,DSN,VOL,RECORDS,
                             INSERTS,DELETES,EXCPS,%CI,%CA,
                             BUFSIZE,CISIZE,CICA)
              SORT     FIELD=(CLUSTER,DSN)
              PRINT    ENABLE=AUTOSTACK,SORTALLOC=YES,PAGEWIDTH=80
```

The generated report will look like:

VSAM INFO FOR SELECTED CLUSTERS

CLUSTER NAME		RECORDS			
DATA SET NAME	VOLSER	INSERTS	%CI	BUFSI	CI
-----	-----	DELETES	%CA	CISIZ	CA
-----	-----	EXCPS	-----	-----	-----
USER1.KSDS	ABC123	290	10	9216	150
USER1.KSDS.DATA		340	20	4096	
		25			
		267			
USER1.KSDS	ABC456	1	0	0	31
USER1.KSDS.INDEX		0	0	1024	
		0			
		20			

54.40 CONTINUED . . .

**REPORT ON
ABR BACKUPS**

FDREPORT is to report on the current ABR backup for all data sets with the index of 'PAYROLL'. The page width is expanded to 132 characters. Associated DSCB information is also requested. All online volumes beginning with certain prefixes are to be searched. The report will be sorted by DSN. The summary will show all ABR backup volumes required to restore all of the payroll data sets, which might be useful as a tape pick list. The backup date will be printed with a 2-digit year; note that the default in V5.3 and above is DATEFORMAT=YYYYDDD for YEAR2000 support.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//ABRSUM      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='PAYROLL BACKUP REPORT'
              REPORT   FIELD=(SPLDSN,VOL,LRDATE,DSORG,RECFM,
                          BLKSIZE,LRECL,SIZE,SIZEUSED,BKINFO)
              SELECT   DSG=PAYROLL,VOL=(PAY*,PROD*)
              SUMMARY   FIELD=(BKVOL)
              PRINT    PAGEWIDTH=132,SORT=COMBINE,SORTALLOCATE=YES,
                      DATEFORMAT=YYDDD
```

The generated report will look like:

PAYROLL BACKUP REPORT

DATA SET NAME	VOLSER	...	BKDATE	BKSUFFIX	FILE	TAPE	VOLUME(S)
PAYROLL.FILE1	PAY001	...	96.304	C1028400	4	BV1048	BV1050
PAYROLL.FILE2	PROD23	...	96.308	C1027902	19	BV1056	

The summary will look like:

```
FINAL TOTALS
VALUE SUMMARY OF BKVOL --- TOTAL NUMBER OF VALUES-----17
      BKVOL-----BV1044 (      1) BV1048 (      3) BV1050 (      3)
```

54.40 CONTINUED . . .

EXECUTE PREDEFINED REPORT FDREPORT is to execute a report which has been set up in advance. The name of this report is VTOCREP1, which is a member in the data set USER.REPORT. This member contains a TITLE, REPORT and PRINT statements. The user supplies a SELECT or XSELECT statement to identify the data sets to be included; in this example, all data sets starting with USER and a numeric digit are selected.

Predefined reports are useful for end-users who have little knowledge of FDREPORT; report formats and controls can be predefined by others and executed by end-users. The FDRLIB members can also contain SELECT statements so that they can be entirely self-contained.

```
//CANREPT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//ABRMAP        DD     SYSOUT=*
//FDRLIB        DD     DSN=USER.REPORT,DISP=SHR
//SYSIN         DD     *
XSELECT        DSN=USER+.*
EXECUTE        REPORT=VTOCREP1
```

If VTOCREP1 contains

```
TITLE          LINE='VTOC REPORT 1'
REPORT         FIELD=(SPLDSN,VOL,SIZEINFO)
PRINT          DATATYPE=CATVTOC
```

then those data sets will be selected from the catalog, VTOC information extracted, and the report will look like:

VTOC REPORT 1

DATA SET NAME	VOLSER	ALLOC	FREE	%FR
USER2.JCL.CNTL	TS0123	10	3	30
USER1.ISPF.ISPPROF	LIBR12	3	0	0

REPORT ON POORLY BLOCKED DATA SETS FDREPORT is to report on all data sets which are using less than 70% of the maximum track capacity for the device they reside on. Generally this will be data sets with small block sizes, or block sizes over a half-track (which waste the rest of the track). The tests on block size and size (tracks) are to eliminate data sets for which no meaningful capacity calculation can be done.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//ABRMAP        DD     SYSOUT=*
//SYSIN         DD     *
TITLE          LINE='POORLY BLOCKED DATA SETS'
XSELECT        %CAPUSED.LT.70,SIZE.GT.0,BLKSIZE.GT.0
REPORT         FIELD=(DSN,VOL,BLKSIZE,DEVTYPE,
                     BLKSTRK,BYTESTRK,TRACKCAP,%CAPUSED)
SUMMARY        FIELD=(DSN,BLKSIZE)
PRINT          ENABLE=ONLINE,DISABLE=INFOMSG
```

The generated report will look like:

POORLY BLOCKED DATA SETS

DATA SET NAME	VOLSER	BLKSZ	DEVTYPE	BLKS TRK	BYTTK	TRCAP	%CP
SYS1.PARMLIB	SYSRES	80	3380	83	6640	47476	14
XYZ.JCL.CNTL	TS0123	800	3380	36	28800	47476	61
YZX.SEQ.DS	PROD01	28672	3380	1	28672	47476	61

The summary will report by volume, and total for all volume in the form:

```
DSN-----25
VALUE SUMMARY OF BLKSIZE --- TOTAL NUMBER OF VALUES-----3
BLKSIZE--- 80 (    1) 160 (    17) 25000 (    7)
```

54.40 CONTINUED . . .

SELECT BY FILTER FDREPORT is to select any online data set that contains a 3-character first level qualifier that starts with the letter 'T' and contains the character string 'YM' anywhere within any qualifier other than the first. Only data sets between 100 and 200 tracks in size will be included. The report will be sorted by dsname within volume, and will be in standard ABR VTOC format (See the PRINT VTOC statement in [Section 53](#)).

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              XSELECT XDSN=T//.*YM*,SIZE.GE.100,SIZE.LE.200
              PRINT   ENABLE=ONLINE,RPTYPE=ABRVTOC,SORT=YES,SORTALLOC=YES
```

The generated report will look like:

ABR VTOC LIST OF VOLUME SERIAL NUMBER TEST01 -- DEVICE TYPE 3380

DATA SET NAME	VOLUME SERIAL	LS CY	VOL SEQ	LAST REF DATE	ABR IND	EXTENT CT	DESCRIPTORS CCC-HH	CCC-HH
TST.NYMASTER	TEST01	00	001	1996.150		1	0296-00	0308-14
TU1.ONLYM.ZEBRA	TEST01	02	001	1996.320	U	2	1250-00	1252-14
							0303-00	0303-14

REPORT ON IAM FILES FDREPORT is to select all online IAM data sets and print statistics about them (see Table 7 in [Section 54.31](#)).

IAM is a product from Innovation which provides a high-performance data-compressed transparent alternative for many VSAM clusters. Contact Innovation for more information.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE   LINE='IAM STATISTICS'
              XSELECT  DSORG=IAM
              REPORT   FIELD=(DSN,VOL,SIZE,IAMUSE)
              PRINT    ENABLE=ONLINE,ENABLE=IAM
```

The generated report will look like:

IAM STATISTICS

DATA SET NAME	VOLSER	ALLOC	RECORDS	OVERFLOW	OVERUSED	%OF
POLICY.MASTER	PROD23	2500	42155	2000	23	2

54.40 CONTINUED . . .

REPORT ON ENHANCED IAM FILES FDREPORT is to select all enhanced IAM data sets on production volumes and print information about them (see Table 2 in [Section 54.31](#)).

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='ENHANCED IAM DATA SETS'
              XSELECT DSORG=IAM,IAMINDIC=ENHANCED
              REPORT  FIELD=(DSN,VOL,SIZE,HIALORBA,HIUSERBA)
              PRINT   ENABLE=(ONLINE,IAM)
```

IDENTIFY FULL PARTITIONED FILE FDREPORT is to locate all partitioned (PO) data sets on TSO volumes and report those that have less than 5 percent free space in the data set or the directory.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='ALMOST FULL TSO PDS DATA SETS'
              XSELECT DSORG=(PO),VOL=TSO*,%FREE<5
              XSELECT DSORG=(PO),VOL=TSO*,%DIRFREE<5
              REPORT  FIELD=(SPLDSN,VOL,SIZEINFO,DIRBINFO)
              PRINT
```

The generated report will look like:

ALMOST FULL TSO PDS DATA SETS

DATA SET NAME	VOLSER	ALLOC	FREE	%FR	DIRBK	DIRFB	%DF
-----	-----	-----	-----	-----	-----	-----	-----
USER1.JCL.CNTL	TS0123	30	1	3	25	23	92
USER1.ISPF.PROF	TS0123	5	2	40	25	1	4

54.40 CONTINUED . . .

REPORT SAME FDREPORT is to report on data sets on a set of disk volumes, once for all data sets sorted by name
DATA TWO within volume, and once for only ICF VSAM sorted by cluster name for all volumes together,
WAYS reporting different fields in each report. To reduce overhead by reading the VTOCs only once, RPTYPE=DATA is used to generate a file of FDREPORT internal records on SYSUT2, then that data is read twice to produce the two reports.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSUT2      DD     UNIT=SYSDA,SPACE=(CYL,(20,5),RLSE)
//SYSIN       DD     *
      DEFAULT      SORTALLOC=YES
*   SELECT ALL DATA WITHOUT REPORT
      SELECT       VOL=PROD*
      PRINT        RPTYPE=DATA
*   GENERATE REPORT 1
      TITLE        LINE='LISTING OF PRODUCTION VOLUMES'
      REPORT       FIELD=(DSN,DSORG,DATES,SIZE,EXTENTS)
      PRINT        DATATYPE=EXTRACT,SORT=YES
*   GENERATE REPORT 2
      CANCEL       SELECT
      TITLE        LINE='LISTING OF PRODUCTION VSAM'
      SELECT       DSORG=EF
      REPORT       FIELD=(CLUSTER,DSN,VOL,DATES,SIZEINFO)
      SORT         FIELD=(CLUSTER)
      PRINT        DATATYPE=EXTRACT
```


54.40 CONTINUED . . .

**REPORT
POORLY
ORGANIZED
VSAM
CLUSTERS**

FDREPORT is to report on all ICF VSAM clusters over 3 Megabytes in size with more than 5 CA splits or more than 20 CI splits, and all clusters with over 16 extents, since all these might be candidates for reorganization. Clusters with the highest CA splits will be listed first.

Note: FDRREORG, a separately priced component, may be used to automate the reorganization of such clusters as well as IAM files and PDSs. [See Section 30](#) for more information.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              DEFAULT SORTALLOC=YES, BYTEFORMAT=MEGABYTES
              TITLE   LINE='VSAM CLUSTERS IN NEED OF REORG'
              XSELECT DSORG=EF, BYTES>3M, CASPLIT>5
              XSELECT DSORG=EF, BYTES>3M, CISPLIT>20
              XSELECT DSORG=EF, NOEXTENT>16
              SORT     FIELD=(CASPLIT, CISPLIT), SEQUENCE=(D, D)
              REPORT   FIELD=(CLUSTER, VOL, PRIALLOC, BYTES,
                             BYTESUSE, CASPLIT, CISPLIT, NOEXTENT)
              PRINT    ENABLE=ONLINE
```

The generated report will look like:

VSAM CLUSTERS IN NEED OF REORG

CLUSTER NAME	VOLSER	PRALO	MBYT	MBUS	CASPL	CISPL	EXT
MASTER.FILE	PROD12	32	22	22	42	175	5
PERM.HISTORY	PROD99	15	125	120	3	27	25

**REPORT DATA
SETS LIKELY
TO GET Sx37**

FDREPORT is to identify data sets likely to get Sx37 (out of space) ABENDs or the equivalent VSAM error. All data sets with less than 10% free space are reported if they have no secondary allocation, or if they have 13 or more extents (50 or more for VSAM).

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE   LINE='DATA SETS WHICH MAY RUN OUT OF SPACE'
              XSELECT SECALLOC.EQ.0,%FREE.LT.10
              XSELECT DSORG.NE.EF,NOEXTENT.GE.13,%FREE.LT.10
              XSELECT DSORG.EQ.EF,NOEXTENT.GE.50,%FREE.LT.10
              REPORT   FIELD=(DSN,DSORG,SECALLOC,%FREE,NOEXTENT,SIZE)
              PRINT    ENABLE=ONLINE
```

The generated report will look like:

DATA SETS WHICH MAY RUN OUT OF SPACE

DATA SET NAME	D/S ORG	VOLSER	SCALO	%FR	EXT	ALLOC
MASTER.FILE	EF	PROD12	10	7	72	780
PROD.CNTL.CARDS	PO	SYS123	0	3	1	20

54.40 CONTINUED . . .

**DIAGNOSE
PROBLEMS IN
VVDS**

FDREPORT is to check the VVDS on specified volumes, reporting on all duplicate records (more than one VVR or NVR for the same data set or component) and orphan records (a VVR or NVR with no corresponding DSCB in the VTOC). Such errors may cause failures when accessing these data sets. Although the DIAGNOSE function of IDCAMS can do the same checks, there is no easy way to do multiple volumes; also, FDREPORT is faster than IDCAMS.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSIN       DD     *
              SELECT  VOL=(TST*,PROD*,SMS*)
              PRINT   RPTYPE=NONE,ENABLE=DIAGNOSEVVDS
```

When errors are found, FDREPORT will print:

```
ORPHANED VVR ON VOLUME TST089 AT RBA 00002511 WITHIN VVDS SYS1.VVDS.VTST089
<dump format print of the VVR>
  REMOVE ORPHANED VVR WITH:
    //STEPXX EXEC PGM=IDCAMS
    //SYSPRINT DD SYSOUT=*
    //DD1 DD UNIT=SYSALLDA,DISP=SHR,VOL=SER=TST089
    //SYSIN DD *
    DELETE TST.VSAM.TEST3.INDEX -
      FILE(DD1) VVR CATALOG(CATALOG.TSOUSER)
    //
FDR421 ICF LOCATE ERROR -- MULTIPLE VVRS (02) ON VOLUME TST089 -- DSN=TST.VSAM.TEST1.DATA
```

54.41 FDREPORT ARCHIVE EXAMPLES

This section shows examples which read an ARCHIVE Control File (DATATYPE=ARCHIVE).

A sample of the generated report is shown after each example; because of space limitations it may be condensed. In examples where SORTing is required, SORTALLOC=YES has been specified to dynamically allocate required SORT files; in your installation you may have to provide SORT JCL.

REPORT ON EXPIRING DATA SETS FDREPORT is to report on all data sets in the archive control file which will expire within the next 30 days, showing information about the location of both archive copies if both exist. The archive date will be printed with a 2-digit year; note that the default in V5.3 and above is DATEFORMAT=YYYYDDD for YEAR2000 support.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//ARCHIVE     DD     DSN=FDRABR.ARCHIVE,DISP=SHR
//SYSIN       DD     *
              DEFAULT DATEFORMAT=YYDDD
              TITLE    LINE='ARCHIVED FILES EXPIRING WITHIN 30 DAYS'
              XSELECT   BKEXDAYS<31
              REPORT    FIELD=(DSN,VOL,BKEXDATE,BKINFO)
              PRINT     DATATYPE=ARCHIVE,COPY=BOTH,SORT=YES,SORTALLOC=YES
```

The generated report will look like:

ARCHIVED FILES EXPIRING WITHIN 30 DAYS

DATA SET NAME	VOLSER	BKEDAT	BKDATE	BKSUFFIX	BKUP FILE	TAPE	VOLUME(S)
USER1.FILE1	TS0001	96.304	95.304	B191304A	2	BA1234	
		96.304	95.304	B291304A	5	BA5432	
USER2.JCL.CNTL	TS0023	96.308	95.308	B191308B	12	BA1234	

SELECTIVE ARCHIVE REPORT FDREPORT is to report on all ARCHIVED data sets which have one of several hi-level indexes, summarized by index. A standard archive report is to be printed. The ARCHIVE Control File whose name is in the ABR option table is to be dynamically allocated and processed. See PRINT ARCHIVE in [Section 53](#) for a sample of the report format.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              DEFAULT SORT=COMBINE,SUM=INDEX,COPY=BOTH
              SELECT   DSG=(FINANCE.,PAYROLL.,LEDGER.)
              PRINT     DATATYPE=ARCHIVE,RPTYPE=ARCHIVE,SORTALLOC=YES
```

54.41 CONTINUED . . .

REPORT ON FDREPORT is to report on all ARCHIVED data sets which were ARCHIVED within the last 2 weeks
SHORT-TERM but which have already been recalled. This might be used to identify data sets which should not
DATA SETS have been ARCHIVED in the first place.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='ARCHIVED DATA SETS RECALLED WITHIN 14 DAYS'
              DEFAULT  DATEFORMAT=MMDDYYYY
              XSELECT  ARCFLAG1=RESTORED,BKDAY<15
              REPORT   FIELD=(DSN,VOL,DSORG,SIZE,BKDATE,ARCFLAGS)
              PRINT    DATATYPE=ARCHIVE
```

The generated report will look like:

ARCHIVED DATA SETS RECALLED WITHIN 14 DAYS

DATA SET NAME	VOLSER	D/S ORG	ALLOC	BKDATE	ARFL1	ARFL2
USER1.FILE1	TS0001	PS	22	10/20/1996	R	A
USER2.JCL.CNTL	TS0023	PS	150	10/22/1996	R	

54.42 FDREPORT CATALOG EXAMPLE

This section shows examples which read the system catalogs. FDREPORT can be directed to gather more information about the cataloged data sets from the VTOCs of the volumes in the catalog (DATATYPE=CATVTOC) or from the ARCHIVE Control File (DATATYPE=CATARCH).

You may also report strictly on the catalog information (DATATYPE=CATALOG) but in this case a very limited set of fields are available.

A sample of the generated report is shown after each example; because of space limitations it may be condensed. In examples where SORTing is required, SORTALLOC=YES has been specified to dynamically allocate required SORT files; in your installation you may have to provide SORT JCL.

**IDENTIFY
MULTI-
VOLUME
DATA SETS**

Identify all data sets which are cataloged to more than one volume. Because of the way that FDREPORT reads the catalogs, the report is naturally sorted by data set name.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='MULTI-VOLUME DATA SETS'
              XSELECT  CATVOLCT>1
              REPORT   FIELD=(DSN,VOL,CATVOLCT)
              PRINT    DATATYPE=CATALOG
```

The generated report will look like:

MULTI-VOLUME DATA SETS

DATA SET NAME	VOLSER	CVC
PROD.MV.FILE	PROD11	2
PROD.MV.FILE	PROD22	2

**REPORT
RECENT TEST
DATA SETS**

FDREPORT is to report on the size of the non-VSAM data sets that have the character string 'TEST' anywhere within their dsname, and that were created within the last two weeks. The data set names will be selected from the system catalogs, then the volumes from the catalog will be accessed to get the rest of the information. The tracks allocated, %FREE, data set name and volume serial are to be printed, sorted by size in descending sequence.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='RECENTLY CREATED TEST DATA SETS'
              REPORT   FIELD=(SIZE,%FREE,DSN,VOL)
              SORT     FIELD=(SIZE),SEQUENCE=(D)
              XSELECT  XDSN=**TEST**,CRDAYS.LE.14,DSORG.NE.EF,CATALOG=YES
              PRINT    SORTALLOC=YES,DATATYPE=CATVTOC
```

The generated report will look like:

RECENTLY CREATED TEST DATA SETS

ALLOC	%FR	DATA SET NAME	VOLSER
450	95	USER1.TEST.JCL	TS0001
295	12	TEST.GL.MASTER1	TEST12

54.43 FDREPORT SMS EXAMPLES

FDREPORT has many uses in a system with SMS (System Managed Storage) active. These are some examples of ways it can be used to aid in the management of an SMS system.

**REPORT BY
MANAGE-
MENT CLASS**

FDREPORT is to report on all online SMS-managed data sets whose SMS management class is TSO1. The data set name, volume, and the SMS class names are to be reported.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='MANAGEMENT CLASS TSO1'
              XSELECT MGMTCLAS=TSO1
              REPORT  FIELD=(DSN,VOL,SMSCLASS)
              PRINT   ENABLE=ONLINE
```

The generated report will look like:

MANAGEMENT CLASS TSO1

DATA SET NAME	VOLSER	STORCLAS	DATACLAS	MGMTCLAS
USER1.TEST.JCL	SMS002	TEST		TSO1
USER3.ISPF.PROF	TSO123	TSO	PROFILE	TSO1

**SUMMARIZE
BY STORAGE
CLASS**

FDREPORT is to scan on all online SMS-managed data sets (STORCLAS.NE.' ' will select SMS data sets since they must all have an assigned storage class), and produce only summaries showing all storage class names in use, and the number of data sets and the number of tracks allocated to those data sets in each class.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRSUM      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='SMS STORAGE CLASS SUMMARY'
              XSELECT STORCLAS.NE.' '
              SORT    FIELD=(STORCLAS),BREAK=(Y)
              SUMMARY FIELD=(STORCLAS,DSN,SIZE)
              PRINT   ENABLE=ONLINE,RPTYPE=NONE,
                      DISABLE=INFOMSG,SORTALLOC=YES
```

The summary will look like:

```
SUBTOTAL -- STORCLAS--GS812
DSN----2    SIZE---75
SUBTOTAL -- STORCLAS--PRODDB
DSN----27   SIZE--4250
```

54.43 CONTINUED . . .

REPORT BY STORAGE GROUP FDREPORT is to report on all data sets in two specific SMS storage groups. For PDS (PO) and PDSE (POE) data sets, member counts are displayed.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='SMS STORAGE GROUPS PROD AND TEST'
              XSELECT STORGRP=(PROD,TEST)
              SORT    FIELD=(STORGRP,DSN)
              REPORT  FIELD=(SPLDSN,STORGRP,VOL,DSORG,LRDATE,SIZE,MEMBERS)
              PRINT   SORTALLOC=YES
```

The generated report will look like:

SMS STORAGE GROUPS PROD AND TEST

DATA SET NAME	STORGRP	VOLSER	D/S ORG	LRDATE	ALLOC	MEMBER
PAYROLL.MASTER	PROD	PROD02	PS	1996.305	750	0
PROD.CARD.CNTL	PROD	PROD12	POE	1996.301	30	27

PREPARE FOR SMS CONVERSION FDREPORT is report on data sets that are ineligible for SMS conversion, so that they can be manually handled. Ineligible data sets are ISAM, non-ICF VSAM, unmovable and uncataloged data sets. CATALOG=NO selects uncataloged data sets, while CATALOG=ERR selects those cataloged to another volume. In the report, the catalog status is printed, as well as the volume to which the data set is cataloged, if any.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='DATA SETS INELIGIBLE FOR SMS'
              XSELECT DSORG=(U,IS,AM),DSORG.NE.EF,VOLG=TSO
              XSELECT CATALOG=NO,VOLG=TSO
              XSELECT CATALOG=ERR,VOLG=TSO
              REPORT  FIELD=(DEFAULTS,CATALOG,CATVOL)
              PRINT
```

54.44 FDREPORT VOLUME EXAMPLES

This section shows examples which produce volume-level reports (DATATYPE=VOLDATA). Available fields are in Table 8 in [Section 54.31](#).

VOLUME USAGE REPORT FDREPORT is to report on the percentage of the volume allocated, plus percentage used for the VTOC, VTOCIX, and VVDS, for all PROD volumes.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='PRODUCTION VOLUME USAGE'
              SELECT   VOL=PROD*
              SORT      FIELD=(VLVOLSER)
              REPORT    FIELD=(VLVOLSER,VLUNIT,VLDEVTYP,VL%UTRKS,
                             VL%UDSCB,VL%UINDX,VL%UVVDS)
              PRINT     SORTALLOD=YES,DATATYPE=VOLDATA
```

The generated report will look like:

PRODUCTION VOLUME USAGE

VOLSER	UAD	DEVTYPE	%TU	%DU	%IU	%VU
PROD07	143	3380	37	42	12	45
PROD12	148	3380	76	17	6	15

VOLUME STATUS REPORT FDREPORT is to report on the mount, usage, and SMS status of every online volume.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//SYSIN       DD     *
              TITLE   LINE='VOLUME STATUS REPORT'
              REPORT    FIELD=(VLVOLSER,VLUNIT,VLDEVTYP,
                             VLUSEATR,VLMOUSTA,VLSMSTAT,VLSMSVST)
              PRINT     ENABLE=ONLINE,DATATYPE=VOLDATA
```

The generated report will look like:

VOLUME STATUS REPORT

VOLSER	UADR	DEVTYPE	USEATTR	MOUNT	SMS STA	VOLUM STATUS
TS0123	0252	3380	STORAGE	RESIDENT	NONE	NONE
DBLRG1	3175	3390	PRIVATE	RESIDENT	MANAGED	ENABLED

54.44 CONTINUED . . .

IDENTIFY DISABLED INDEXED VTOCS FDREPORT is to identify volumes which have an Indexed VTOC (VTOCIX) which has been disabled. A status of YES indicates that the VTOCIX exists but is not active.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT     DD     SYSOUT=*
//ABRMAP       DD     SYSOUT=*
//SYSIN        DD     *
               TITLE   LINE='DISABLED VTOC INDEXES'
               XSELECT  VLINDSTA=YES
               REPORT   FIELD=(VLVOLSER,VLINDSTA)
               PRINT    ENABLE=ONLINE,DATATYPE=VOLDATA
```

The generated report will look like:

DISABLED VTOC INDEXES

```
VOLSER    VTOCIX
-----
TS0123     YES
DBLRG1     YES
```

VOLUME STATISTICS FOR IMPORT FDREPORT is to generate statistics for all online volumes in a tabular format for import into another program. If this is done at regular intervals, the other program might be used to keep a history of the values, generating history, statistical and trend reports. You might also download the report file to a PC for analysis by PC-based programs. The tabular format will contain no titles or page breaks; it will contain one set of headings unless DISABLE=HEADINGS is specified as shown.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT     DD     SYSOUT=*
//ABRMAP       DD     DSN=SYSTEMS.VOLSTAT(+1),UNIT=DISK,
//              SPACE=(TRK,3),DISP=(,CATLG)
//SYSIN        DD     *
               XSELECT  VLINDSTA=YES
               REPORT   FIELD=(VLVOLSER,VLUNIT,VLDEVTYP,VLTRKVOL,VLAOTRK,VLDSCB1)
               PRINT    ENABLE=ONLINE,DATATYPE=VOLDATA,DISABLE=HEADINGS
```

The generated file will look like:

```
TS0123 1A3 3390 50085 25173 236
MVSRES 1C7 3380 39825 37900 1274
```

The program which reads the file must, of course, be aware of the meaning and position of each column.

54.45 FDREPORT PUNCH EXAMPLES

This section shows examples of generating JCL and control cards using FDREPORT data fields (RPTYPE=SELPCH).

**GENERATE
ABR
STATEMENTS**

FDREPORT is to be used as a front-end filter for ABR, selecting data sets to be scratched by SUPERSCRATCH. All data sets whose last index level begins with "LIST" or "TEMP" will be scratched, if they were created more than 1 day ago. FDREPORT will generate ABR control statements in the default format:

SELECT DSN=dsname ,VOL=volume
and write them to the SYSPUNCH temporary data set, which is read by the following ABR step.

```
//SELECT      EXEC   PGM=FDREPORT ,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPUNCH    DD     DSN=&&ABRIN ,UNIT=SYSDA ,SPACE=(TRK,(2,1)),
//            DD     DISP=(NEW,PASS) ,DCB=BLKSIZE=6160
//SYSIN       DD     *
              XSELECT   XDSN=(** .LIST*,** .TEMP*),CRDAYS>1
              PRINT     ENABLE=ONLINE ,RPTYPE=SELPCH
//SUPERSCR    EXEC   PGM=FDRABR ,COND=(0,NE,SELECT)
//SYSPRINT    DD     SYSOUT=*
//SYSPRINT1   DD     SYSOUT=*
//TAPE1       DD     DUMMY
//SYSIN       DD     *
              DUMP      TYPE=SCR ,ONLVOL ,DSNENQ=USE ,MAXCARDS=1000
//            DD     DSN=&&ABRIN ,DISP=(OLD,DELETE)
```

**GENERATE
ABR RESTORE
SELECTS**

FDREPORT is to generate a data set containing SELECT statements specifying the dsn, volume, and ABR generation and cycle for each selected data set, using a punch mask. This might be used at a disaster/recovery site to do data set restores of selected data sets from their most recent ABR backup.

```
//PUNCHMSK    EXEC   PGM=FDREPORT ,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//SYSPUNCH    DD     DSN=SELPCH.OUTPUT ,UNIT=3490 ,DISP=( ,CATLG) ,
//            DD     DCB=BLKSIZE=6160
//SYSIN       DD     *
              XSELECT   XDSN=PAYROLL. **
              PUNCH     FDRLIB=MASK
              PRINT     ENABLE=ONLINE ,RPTYPE=SELPCH
//MASK        DD     *
              SELECT    DSN=<NAME> ,VOL=<VOL> ,
                       GEN=<ABRGEN> ,CYCLE=<ABRCYCLE> ,
                       NVOL=PAY*
```

The generated statements on SYSPUNCH will look like:

```
SELECT DSN=PAYROLL.HOURLY.MASTER ,VOL=PAY001 ,
       GEN=0123 ,CYCLE=005 ,
       NVOL=PAY*
```

54.45 CONTINUED . . .

GENERATE FDREPORT is to generate an IDCAMS jobstream to delete selected clusters and data sets. The
IDCAMS job will be submitted directly to the JES internal reader for execution. The punch mask is read from
DELETE a member of the FDRLIB library.

```
//PUNCHMSK EXEC PGM=FDREPORT,REGION=0M
//SYSPRINT DD SYSOUT=*
//SYSPUNCH DD SYSOUT=(A,INTRDR)
//FDRLIB DD DSN=USER1.FDREPORT,DISP=SHR
//SYSIN DD *
XSELECT XDSN=OLDFILES.*
PUNCH MASKNAME=IDCDEL,ECHO
PRINT DATATYPE=CATALOG,RPTYPE=SELPCH
```

Member IDCDEL of library USER1.FDREPORT contains:

```
)PREFIX
//DELETE JOB (ACCT),CLASS=M,MSGCLASS=X
//* DELETE OLD FILES
//DELETE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
)ENDPREFIX
DELETE <NAME> PURGE
```

The job submitted to the internal reader will look like:

```
//DELETE JOB (ACCT),CLASS=M,MSGCLASS=X
//* DELETE OLD FILES
//DELETE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE OLDFILES.JCL.CNTL PURGE
DELETE OLDFILES.RECORDS.KSDS PURGE
```

CONTROL ABR Here is an innovative use of FDREPORT PUNCH processing to circumvent a limitation of ABR. In
ORDER OF ABR it is convenient to use MOUNT statements or VOL= or VOLG= operands on SELECT
DISK statements to define the volumes which ABR is to process in a given backup step. Since you can
PROCESSING specify volume groups and SMS storage groups for selection, new volumes and deleted volume
are automatically picked up.

However, ABR will process these volumes in the order of their UCBs (Unit Control Blocks) in the MVS system; these are usually in order by device address. If you are doing concurrent backups (multiple TAPEX DD statements), this means that ABR will tend to process several volumes on the same control unit and device string at once. This may overload the data paths to the control unit or disk and degrade performance. The only solution is to use DISKxxxx DD statements in the ABR step, causing ABR to process the volumes in the order of the DD statements, but this requires changes to the ABR job whenever the volume list changes.

To circumvent this, this jobstream will generate DISKxxxx DD statements sorted by the last digit of each volume's device address. In other words, all addresses ending in A will be sorted first, such as 01CA and 13DA, followed by those ending in B through F, and 0 through 9. This will tend to spread the concurrent backups out over several control units and strings. This will also improve backup performance when outputting to high-speed tapes such as IBM Magstar 3590 and STK Redwood.

- Step REPORT uses PUNCH to generate a DISKxxxx DD statement for each selected volume, where "xxxx" is set to the volume's device address, such as DISK01A2.
- Step SORT sorts the generated DD statements on the last digit of the device address (column 10).
- Step SUBMIT uses IEBGENER to submit an ABR job (an incremental backup in this example) with those sorted DD statements appended.

54.45 CONTINUED . . .

This technique may significantly reduce in your ABR elapsed time, but results will differ from installation to installation, depending on the DASD device, control unit and channel configuration.

The example jobstream is shown on the next page. All control statements, including the punch mask and the ABR step JCL, are shown inline, but you may want to move them to control statement libraries and change the DD * and DD DATA statements to point to those members.

```
//REPORT EXEC PGM=FDREPORT
//SYSPRINT DD SYSOUT=*
//SUMPRT DD SYSOUT=*
//PCHJCL DD DSN=JCL,UNIT=VIO,SPACE=(TRK,(5,2)),
// DISP=(,PASS),
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=800)
//SYSIN DD *
*
* FIND ALL SELECTED DISK VOLUMES AND GENERATE A DISKXXXX DD
* STATEMENT FOR EACH.
*
* MODIFY THE FOLLOWING XEXCLUDE AND XSELECT STATEMENTS TO SELECT
* THE VOLUMES YOU REQUIRE AND/OR SMS STORAGE GROUPS YOU REQUIRE
*
XEXCLUDE VOL=IDPWK*
XSELECT VOLG=IDP
XSELECT STORGRP=DATABASE
*
* MODIFY THE FOLLOWING EXCLUDE AND SELECT STATEMENTS TO SELECT
* THE VOLUMES YOU REQUIRE OR SMS STORAGE GROUPS YOU REQUIRE
*
PUNCH FDRLIB=MASKDD,ECHO
PRINT RPTYPE=SELPCH,DATATYPE=VOLDATA,PCHDDNAME=PCHJCL
*
* THE FOLLOWING PUNCH MASK WILL GENERATE A DISKXXXX DD STATEMENT
* FOR EACH VOLUME WITH XXXX REPLACED BY ITS 4-DIGIT DEVICE ADDRESS.
*
//MASKDD DD DATA
//DISK<UNIT> DD DISP=SHR,UNIT=<UNIT>,VOL=SER=<VOL>
/*
//SORT EXEC PGM=SORT
/*
/* THIS SORT WILL SORT THE DISKXXXX DD STATEMENTS BY THE LAST
/* DIGIT OF THEIR DEVICE ADDRESS
/*
//SYSOUT DD SYSOUT=*
/*
/* IF YOUR SORT PRODUCT REQUIRES ADDITIONAL DD STATEMENTS,
/* ADD THEM HERE
/*
//SORTIN DD DSN=JCL,DISP=(OLD,DELETE)
//SORTOUT DD DSN=JCL2,UNIT=VIO,SPACE=(TRK,(5,2)),
// DISP=(,PASS),
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=800)
//SYSIN DD *
SORT FIELDS=(10,1,CH,A)
//SUBMIT EXEC PGM=IEBGNER
/*
/* SUBMIT THE ABR JOB SHOWN BELOW (AFTER THE SYSUT1 DD) AND
/* APPEND THE SORTED DISKXXXX DD STATEMENTS
/*
//SYSPRINT DD SYSOUT=*
```

54.45 CONTINUED . . .

```
//SYSIN      DD DUMMY
//SYSUT2     DD SYSOUT=(*)
//SYSUT1     DD DATA
//ABRJOB     JOB . . . (change for your installation's requirements)
//ABRBKUP    EXEC PGM=FDRABR,REGION=0M
//SYSPRINT   DD SYSOUT=*
//SYSPRIN1   DD SYSOUT=*
//SYSPRIN2   DD SYSOUT=*
//SYSPRIN3   DD SYSOUT=*
//SYSUDUMP   DD SYSOUT=*
//TAPE1      DD DSN=ABR1,UNIT=CART,DISP=(,KEEP),VOL=(,,,255)
//TAPE2      DD DSN=ABR1,UNIT=CART,DISP=(,KEEP),VOL=(,,,255)
//TAPE3      DD DSN=ABR1,UNIT=CART,DISP=(,KEEP),VOL=(,,,255)
             DUMP TYPE=ABR
/*
//           DD DSN=&&JCL2,DISP=(OLD,DELETE)
```

54.46 FDREPORT TAPE EXAMPLES

NOTE: the PRINT TVTOC function of program FDRABRP (see Section 53) can also be used to generate fixed-format reports from FDR-format backup tapes. PRINT TVTOC is available even to FDR customers who are not also ABR customers, while FDREPORT is only available if you have licensed ABR or FDREPORT.

MAP BACKUP TAPE FDREPORT is to read a FDR-format backup tape and report on the data sets included in that backup using the default report. The JCL shows an ABR full-volume backup, but it could be any backup created by FDR, DSF, ABR, or SAR.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//TAPE001     DD     DSN=FDRABR.VPROD01.C1002300,DISP=SHR
//SYSIN       DD     *
               TITLE   LINE='DATA SETS ON BACKUP <TVTOCDSN>'
               PRINT    DATATYPE=TVTOC
```

The generated report will look like:

DATA SETS ON BACKUP FDRABR.VPROD01.C1002300

		D/S						
DATA SET NAME	VOLSER	ORG	RECFM	BLKSZ	LRECL	ALLOC	FREE	%FR
PAYROLL.TRANS	PROD01	PS	FB	3120	80	10	0	0
PROD.JCL.CNTL	PROD01	PO	FB	6160	80	5	2	40

MAP BACKUP TAPES FDREPORT is to read several FDR backup tapes and report on the PDS data sets included in those backups. The ENABLE=TAPERREAD option is included so that FDREPORT will read the entire backup to extract PDS directory information. The report is sorted by backup data set name and by size within backup.

```
//REPORT      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT    DD     SYSOUT=*
//ABRMAP      DD     SYSOUT=*
//TAPE1       DD     DSN=PROD.BACKUP.VTS0001(0),DISP=SHR
//TAPE2       DD     DSN=PROD.BACKUP.VTS0002(0),DISP=SHR
//TAPE3       DD     DSN=PROD.BACKUP.VTS0003(0),DISP=SHR
//SYSIN       DD     *
               TITLE   LINE='PO DATA SETS ON BACKUP <TVTOCDSN>'
               SELECT   DSORG=PO
               SORT      FIELD=(TVTOCDSN,SIZE),SEQUENCE=(A,D),
                           BREAK=(YES,NO)
               REPORT    FIELD=(SPLDSN,VOL,SIZE,MEMBERS)
               PRINT     DATATYPE=TVTOC,ENABLE=TAPERREAD,SORTALLOC=YES
```

The generated report will look like:

PO DATA SETS ON BACKUP PROD.BACKUP.VTS0001.G0123V00

DATA SET NAME	VOLSER	ALLOC	MEMBER
USER1.JCL.CNTL	TS0001	20	125

54.46 CONTINUED . . .

PRINT The ARCHIVE Control File contains only limited information about the original disk data sets
ARCHIVE INFO recorded in it. This jobstream will select a data set from the control file and submit a second job
 which will read the ARCHIVE backup that contains it and report on selected fields.

```
//REPORT1      EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//SYSPUNCH      DD     SYSOUT=(A,INTRDR)    <== SUBMIT TO JES
//SYSIN         DD     *
                SELECT  DSN=datasetname      <== specify data set name here
                PUNCH   FDRLIB=MASK,ECHO
                PRINT    DATATYPE=ARCHIVE,RPTYPE=SELPCH,COPY=1
//MASK          DD     DATA,DLM=$$
//jobname       JOB    required parameters
//REPORT2       EXEC   PGM=FDREPORT,REGION=0M
//SYSPRINT      DD     SYSOUT=*
//TAPE1         DD     UNIT=<BKDEVTYPE>,DISP=SHR,LABEL=<BKFILENO>,
//              VOL=SER=<BKVOL>,DSN=FDRABR.V<VOL>.<BSUFFIX>
//SYSIN         DD     *
                TITLE   LINE='ARCHIVED DATA SET <NAME>'
                SELECT  DSN=<NAME>
                REPORT  FIELD=(SPLNAME,VOL,DATES,SIZEINFO,EXTENTS)
                PRINT    DATATYPE=TVTOC
$$
```

54.47 FDREPORT YEAR 2000 EXAMPLES

These examples demonstrate the support in FDREPORT for 4-digit years, including the year 2000 and beyond.

**ARCHIVES
EXPIRING
BEYOND 1999**

FDREPORT is to report on archived data sets whose archive will expire in the next century. Note that the date on the XSELECT statement can contain a period for readability (2000.001) or the period can be omitted (2000001). The DATEFORMAT= operand affects only the format of the printed report. Dates on (X)SELECT statements can be entered with either 2- or 4-digit years (2-digit years are assumed to be 19xx).

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='ARCHIVED DATA SETS EXPIRING BEYOND 1999'
              DEFAULT DATEFORMAT=YYYYDD
              XSELECT  BKEXDATE,GE,2000.001
              REPORT   FIELD=(DSN,VOL,DSORG,SIZE,BKDATE,BKEXDATE)
              PRINT    DATATYPE=ARCHIVE
```

The generated report will look like:

ARCHIVED DATA SETS EXPIRING BEYOND 1999

		D/S			
DATA SET NAME	VOLSER	ORG	ALLOC	BKDATE	BKEXDATE
USER1.FILE1	TS0001	PS	22	1998.056	2001.056
USER2.JCL.CNTL	TS0023	PO	150	1997.279	2000.279

**IDENTIFY
DATA SETS
WITH
PERMANENT
RETENTION**

FDREPORT is to report on all DASD data sets with expiration dates of 99365 or 99366 (both are considered permanent "never scratch" dates by MVS). Although current releases of MVS will bump any expiration calculated from a RETPD= value from 1999.365 to 2000.001, it is possible that a data set created under an earlier release of MVS might have a calculated date which is now incorrectly considered permanent. This report can be reviewed to ensure that all such data sets should really be considered permanent.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='PERMANENT RETENTION DATA SETS'
              XSELECT EXPDATE=(99365,99366)
              REPORT  FIELD=(SPLDSN,VOL,DATES)
              PRINT   ENABLE=ONLINE,DATEFORMAT=MMDDYYYY
```

The generated report will look like:

PERMANENT RETENTION DATA SETS

DATA SET NAME	VOLSER	CRDATE	LRDATE	EXDATE
USER1.JCL.CNTL	TS0001	05/07/1997	06/22/1997	12/31/1999

54.48 FDREPORT HSM EXAMPLES

Please review [Section 54.32](#) for details on the FDREPORT field names which are available when you are reporting on data in a DFSMSHsm or DFHSM MCDS or BCDS data base.

MCDS REPORT FDREPORT is to report on data from the HSM MCDS. Only certain data sets are selected, and certain fields which are valid for the MCDS are included in the report. The MCDS will be dynamically allocated.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//SYSIN       DD    *
              TITLE  LINE='HSM MIGRATED DATA SETS'
              DEFAULT MCDSCLUSTER=HSM.MCDS
              XSELECT XDSN=USER1**
              REPORT  FIELD=(DSN,VOL,DSORG,SIZE,SIZEFREE,ADATE,ATIME,ADAYS)
              PRINT   DATATYPE=MCDS
```

The generated report will look like:

HSM MIGRATED DATA SETS

	D/S						
DATA SET NAME	VOLSER	ORG	ALLOC	FREE	ADATE	ATIME	ADAYS
-----	-----	-----	-----	-----	-----	-----	-----
USER1.TEST.ESDS	SMS802	EF	1	0	1995.116	11.54.49	1055

BCDS REPORT FDREPORT is to report on data from the HSM BCDS. Only certain data sets are selected, and certain fields which are valid for the BCDS are included in the report. The BCDS name is specified in JCL.

```
//REPORT      EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//BCDSDD      DD    DISP=SHR,DSN=HSM.BCDS
//SYSIN       DD    *
              TITLE  LINE='HSM BACKUPS OF DATA SETS'
              XSELECT XDSN=USER1**
              REPORT  FIELD=(DSN,VOL,DSORG,SIZE,ADATE,ATIME,ADAYS)
              PRINT   DATATYPE=BCDS
```

The generated report will look like:

HSM BACKUP OF DATA SETS

	D/S						
DATA SET NAME	VOLSER	ORG	ALLOC	ADATE	ATIME	ADAYS	
-----	-----	-----	-----	-----	-----	-----	-----
USER1.AC.DATA	SMS802	PS	5	1998.048	20.15.35	16	

54.50 FDRSRS -- SEARCH, REPORT AND SERVICES DIALOG

INTRODUCTION The FDRSRS ISPF dialog provides a fast and easy way of selecting, reporting and performing services against data sets and volumes. FDRSRS uses FDREPORT to do most of the data gathering and formatting, but it makes the selection and display of the data very easy to do. Various commands and functions can be executed against the data displayed. It can be used by all types of ISPF users, from DASD managers to end-users. FDRSRS is divided into a Data Set Application and a Volume Application.

SRS is superior to the data set and volume functions of ISMF and ISPF 3.4 in speed, flexibility, and ease of use.

DATA SETS The Data Set Application selects data sets from a variety of sources, reports the requested data set attributes (VOLSER, DSORG, RECFM, BLKSIZE, etc. over 150 selectable attributes), and performs ABR and other services against the selected data sets. The sources that may be searched are system catalogs, VTOCs of online volumes, the ABR Archive Control file, the ABR Scratch Catalog, or an extract file created by FDREPORT or FDRSRS.

The Data Set Selection Criteria Panel displays rows of data set attributes and columns where the user may specify selection, reporting, summary, and sorting criteria. Additionally, each row has columns describing the attribute.

The Data Set List Panel displays rows containing the selected data sets and columns containing the requested (or defaulted) data set attributes. Additionally, each row contains a command column where the user may request ABR services (such as Recall, Reorg, Copy), ISPF services (such as Edit, Browse), TSO commands (such as DELETE, LISTDS), CLISTs, and REXX execs. The format of the data set list can be easily modified. You can specify which data set attributes should be displayed and in what order (or let SRS pick the order). You can print out hardcopy of any Data Set List.

VOLUMES The Volume Application selects online volumes and reports the requested volume attributes (over 50 selectable attributes).

The Volume Selection Criteria Panel displays rows of volume attributes and columns where the user may specify selection, reporting, summary, and sorting criteria. Additionally, each row has columns describing the volume attribute and its output length.

The Volume List Panel displays rows containing the selected volumes and columns containing the requested (or defaulted) volume attributes.

SAVING LISTS The selection, report, summary and sort criteria specified may be saved on disk as a member in the Selection Criteria set of libraries. To allow for further customization of the Selection Criteria, a user comments area is provided and each row contains a command column where commands may be issued to delete unwanted rows, repeat rows, or move rows before or after other rows. This allows users to save commonly used selection and reporting criteria for reuse. Storage Managers may also setup criteria for use by other users.

A Data Set or Volume List (the output of SRS) may be also saved on disk for later use. When a List is saved, its associated Selection Criteria is also saved. A saved List may be refreshed in the future using its original Selection Criteria. Additionally, the Selection Criteria associated to a Data Set or Volume List may be modified and processed in the future. The Data Set or Volume List is saved as a member in the List set of libraries.

TRY IT! The following sections provide an overview of SRS with simple examples. The power of SRS can be appreciated only by experimenting with it. Context-sensitive HELP is available at every point within SRS. The panels you see may differ slightly from those printed here.

NOTE: Module FDRSRSA must be placed in the TSO authorized program name table before FDRSRS is invoked. Please refer to [Section 90](#) for instructions on updating this table.

54.50 CONTINUED . . .

SAMPLE SELECTION CRITERIA PANEL

-----DATASET SELECTION: DEFAULT -----LINE 1/142 COL 4:6/8
COMMAND ==> _ SCROLL==> HALF

Read Save Submit Find Locate Extract Options Help

FIELD SELECTION VALUE REPORT SORT

SOURCE FIELDS

DSNAME ==> DF.**

VOL ==>

SOURCE ==> CATALOG (Catalog Volume Archive Scratch Extract)

CATALOGN ==>

ARCDSN ==>

VTOC FIELDS

DEFAULTS ==>

UNIT ==>

DEVTYPE ==>

VOLSQ ==>

DATES ==>

CRDATE ==>

CRDAYS ==>

SMS FIELDS

SMSCLASS ==>

STORCLAS ==>

MGMTCLAS ==>

DATACLAS ==>

STORGRP ==>

SMSFLAGS ==>

VVDS FIELDS

CLUSTER ==>

SPLCLS ==>

SCLUSTER ==>

VSAMID ==>

Example of some VTOC, SMS and VSAM fields that can be Selected, Sorted and/or Reported

Data set DF.ABRPANEL is archived. Specifying B (Browse) in the command column will cause an automatic recall of this data set before going into Browse.

Printd command allows the user to change the order in which the data set attributes are printed and the number of copies to print.

-----DATASET LIST: DSLIST -----LINE 1/77 COL 4:10/10
COMMAND ==> _ SCROLL==> HALF

	Read	Refresh	Next	Message	Printd	Help	
COMMAND	VOLSER	DSO	RECFM	BKSIZ	LRECL	ALLOC	FREE
DF.ABRMAP	MIGRAT	ARCHIVED					
DF.ABRPANEL	MIGRAT	ARCHIVED					
DF.ALTLIB.CLIST	SCR083	P0	FB	8000	80	15	14
DF.ANSWERS	IDPLB0	PS	VB	6233	84	1	0

SAMPLE DATA SET LIST PANELS

Data set DF.ARCHIVE is no longer required. Specifying DELETE will cause the data set to be scratched and uncataloged.

-----DATASET LIST: DSLIST -----LINE 1/77 COL 4:10/10
COMMAND ==> _ SCROLL==> HALF

	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help
COMMAND	DATA SET NAME	VOLSER	DSO	RECFM	BKSIZ	LRECL	ALLOC	FREE	
DF.APF.LOAD	IDPLB1	P0	U		32760		134	40	
DELETE DF.ARCHIVE	IDPLB0	DA	F		23392		1	0	
DF.ASM.MACLIB	IDPLB4	P0	FB		3280	80	12	3	
REORG DF.ASM.SOURCE	IDPLB4	P0	FB		19040	80	315	65	
RESTORE DF.CALL.LOG	IDPLB1	P0	FB		6192	72	23	2	

54.50 CONTINUED . . .

The SRS dialog is invoked by selecting option S in the FDR/ABR Primary Options Menu. For faster access, or if you wish to give users access to SRS without the other ABR dialogs, you may add the SRS option to the ISPF system command table. [Section 54.57](#) discusses fastpaths for invoking the SRS dialog, bypassing some of the preliminary SRS panels.

For simplicity, the panel names and options shown in the rest of this section assume that option A on the ISPF main menu is used for accessing the FDR/ABR Primary Options menu. So, you can get to the SRS Primary Menu by entering "S" on the FDR/ABR Primary Options Menu, or "A.S" on the ISPF Primary Options Menu.

**PANEL A.S SRS
PRIMARY
MENU**

```

----- FDRSRS - Primary Menu -----
OPTION  ===> 1                                SCROLL ===> HALF
                                           More:      +

 0  OPTIONS  - Set Dialog Options and Defaults

 1  SELECT   - Data Set Selection
              Name ===> *                      (*, member name, or blanks)

 2  DSLIST   - Data Set List - display saved
              Name ===> *                      (*, member name, or blanks)

 3  SELVOL   - Volume Selection
              Name ===> DEFAULT                (*, member name, or blanks)

 4  VOLLIST   - Volume List - display saved
              Name ===> *                      (*, member name, or blanks)

OR Select one of the following services:
- ARCDL      - DELETE DATA SET ENTRY IN THE ARCHIVE FILE
- ARCHIVE    - ARCHIVE DATA SET
- ARCMOD     - MODIFY DATA SET ENTRY IN THE ARCHIVE FILE
- ARCRECAT   - RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL
- ARCRESET   - RESET DATA SET ENTRY IN THE ARCHIVE FILE
- BACKAPPL   - DATA SET APPLICATION BACKUP
- BACKUP     - ADD BACKUP REQUEST TO REMOTE QUEUE
- COPY       - COPY DATA SETS
- FDRREORG   - FDRREORG DATA SET REORGANIZATION
- MOVE       - MOVE DATA SETS
- REMOVEA    - DELETE ARCHIVE RESTORE REQUEST FROM REMOTE QUEUE
- REMOVEB    - DELETE BACKUP RESTORE REQUEST FROM REMOTE QUEUE
- REORG      - COMPRESS PDS DATA SETS
- RESETARC   - DELETE ARCHIVE REQUEST FROM REMOTE QUEUE
- RESETBKP   - DELETE BACKUP REQUEST FROM REMOTE QUEUE
- RESTAPPL   - RESTORE DATA SETS FROM APPLICATION BACKUP
- RESTARC    - RESTORE DATA SETS FROM ARCHIVE
- RESTBKP    - RESTORE DATA SETS FROM BACKUP
- SIMREORG   - SIMULATE FDRREORG DATA SET REORGANIZATION

```

The SRS primary menu is used to select the SRS function desired. **Option 0** allows each user to set options and defaults that affect his/her use of SRS. **Options 1 and 2** select the data set application, **options 3 and 4** select the volume application. The odd-numbered options actually select data sets or volumes; you can optionally retrieve a saved selection criteria. The even-numbered options retrieve a data set or volume list previously selected and saved.

The various name fields refer to members in SRS libraries of saved selection criteria and lists. The names of these libraries can be specified via option 0.2 or overridden by specifying blanks for the member name. The libraries can include private user SRS libraries and/or shared or installation-wide libraries.

54.50 CONTINUED . . .

The Selection Name field specifies the member name of a previously saved Selection Criteria, or one of the following special values:

- | | |
|---------|---|
| DEFAULT | this is the built-in selection criteria starter set containing all the available fields. |
| Blanks | display a panel containing the names of the libraries that will be used to read the Selection Criteria. |
| * | display a panel containing the names of the Selection Criteria that were previously saved. |

The List Name field specifies the member name of a previously saved List, or one of the following special values:

- | | |
|--------|---|
| Blanks | display a panel containing the names of the libraries that will be used to read the List. |
| * | display a panel containing the names of the Lists that were previously saved. |

NOTE: *When the first SAVE command is specified, if the library specified in the SRS options as the READ/WRITE library for that type of save does not exist, the dialog will display a panel containing allocation parameters for creating the Selection Criteria or List library. There is no need to pre-allocate the FDRSRS libraries.*

The SRS primary menu also presents a list of FDR functions which can be invoked from SRS. Placing a S next to any one of them will invoke the SRS dialog for that function. These are the same functions which can be selected from a SRS display of the results of a data set or volume report; in that case SRS will fill in fields on the following panels with information about the data set or volume selected. Since no search is done when these functions are invoked from this panel, the user may have to fill in additional information.

54.51 SRS DIALOG OPTIONS AND DEFAULTS PANELS

The FDRSRS Defaults and Options Panels are displayed by selecting option 0 from the main FDRSRS panel, or by issuing the command OPTIONS (or O) from the Selection Criteria panel or the List panel.

**PANEL A.S.0
SRS OPTIONS
AND
DEFAULTS**

```

----- FDRSRS - User Options and Defaults -----
OPTION  ==>

Select an option or use PF8/20 (DOWN) to display the options panels in sequence

    1 - Processing Options

    2 - Data Set Names

    3 - Batch JCL submission options

    4 - Report format options

- - - - -

FF - FDR Function Commands (saved in a private table library)

FS - FDR Function Commands (for all users, saved in a common table library)

```

The panels which are invoked by each of these options are displayed on the following pages, to show you the options that are available. However, they are not described here. Please refer to the HELP tutorial for current and detailed information on the user options and defaults.

Options changed on these panels are permanently stored in the user's ISPF profile data set. To change installation-wide default option values, use ISPF to edit member FDRSRSD in the FDR panel library and follow the instructions contained in this member. (Note: the modification of an installation-wide default option value will only affect those users who have not explicitly changed the option value.) For example, you may want to provide libraries of saved queries and lists that any user can invoke; these can be specified as read-only libraries on the Default Data Set Names panel.

Most users of SRS will probably not need to modify the options.

54.51 CONTINUED . . .

PANEL A.S.0.1
SRS
PROCESSING
OPTIONS

```

----- FDRSRS - Processing Options -----
OPTION  ==>

Number of selected data sets after which to suspend selection and display list:
  Suspend ==> 100

Search only the specified catalog (i.e., do not switch to connected catalogs):
  Onecat  ==> YES          (yes | no)

Report errors involving OS CVOs (unsupported) and offline catalogs:
  Caterr  ==> NO          (yes | no)

Select all entries from the catalog (including tape data sets):
  Allent  ==> NO          (yes | no)

Execute the Selection Criteria when the Enter key is pressed, or RUN is entered:
  Execute ==> ENTER       (Enter | Run)

Convert data set name into data set name filter for selection:
  Convert ==> YES         (yes | no | dsg)

```

PANEL A.S.0.2
SRS DEFAULT
DATA SET
NAMES

```

----- FDRSRS - Default Data Set Names -----
OPTION  ==>

Default Selection Criteria library names:
  Read/Write dsn ==> 'FDRABR.SRS.SELECT'
  Read only  dsn ==>
  Read only  dsn ==>

Default Data Set List/Volume List library names:
  Read/Write dsn ==> 'FDRABR.SRS.LIST'
  Read only  dsn ==>
  Read only  dsn ==>

Default FDREPORT Extract data set name:
  Extract dsname ==> 'FDRABR.SRS.EXTRACT'

Default ABR Archive Control File name (or blank):
  Archive dsname ==>

Default Catalog name where to search (or blank to use the master catalog):
  Catalog dsname ==>

```

54.51 CONTINUED . . .

PANEL A.S.0.3
SRS BATCH
JCL OPTIONS

```

----- FDRSRS - Batch Submission Options -----
OPTION  ==>

Job Statement Information:
  ==> //useridA JOB (ACCOUNT),'NAME',NOTIFY=userid
  ==> /**
  ==> /**
  ==> /**

FDR Program Library for STEPLIB DD (blank if LINKLIST) :
  Steplib ==>

SYSOUT Class  ==> *
```

PANEL A.S.0.4
SRS REPORT
FORMAT
OPTIONS

```

----- FDRSRS - Report Format Options -----
OPTION  ==>

Date Format  ==> MMDDYYYY          ( YYDDD | MMDDYY | DDMMYY | YYYYDDD |
                                   MMDDYYYY | DDMMYYYY )
Byte Format  ==> KILOBYTES         ( bytes | kilobytes | megabytes )

Fold Dsname  ==> YES              ( yes | no )

Backup Copy  ==> EITHER           ( both | either | 1 | 2 )

Oldbackup    ==> ALL              ( all | cur | (n,n,n,...) )

Optional FDREPORT DEFAULT statement:
  ==> DEFAULT
```

Note that the default for date format in V5.3 (MMDDYYYY) is changed from the default in earlier releases (YYDDD) for Year2000 support; this default is also different from the default in FDREPORT (YYYYDDD).

54.51 CONTINUED . . .

PANEL
A.S.O.FF/FS
FDR
FUNCTION
COMMANDS

```

----- FDRSRS - FDR Function Command ----- Row 1 of 19
COMMAND ==> SCROLL ==> PAGE

Select one of the following row selection codes, or press PF3 (END) to exit.
  S - Select   R - Repeat   I - Insert   D - Delete

Sel   Command   Command
Code  Name      Description
-----
      ARCDL     DELETE DATA SET ENTRY IN THE ARCHIVE FILE      DEFAULT
      ARCHIVE   ARCHIVE DATA SET                                DEFAULT
      ARCMOD    MODIFY DATA SET ENTRY IN THE ARCHIVE FILE      DEFAULT
      ARCRCAT   RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL    DEFAULT
      ARCRES    RESET DATA SET ENTRY IN THE ARCHIVE FILE      DEFAULT
      BACKAPPL  DATA SET APPLICATION BACKUP                    DEFAULT
      BACKUP    ADD BACKUP REQUEST TO REMOTE QUEUE              DEFAULT
s     COPY      COPY DATA SETS                                  DEFAULT
      FDRREORG  FDRREORG DATA SET REORGANIZATION              DEFAULT
      MOVE      MOVE DATA SETS                                  DEFAULT
      REMOVEA   DELETE ARCHIVE RESTORE REQUEST FROM REMOTE QUEUE  DEFAULT
      REMOVEB   DELETE BACKUP RESTORE REQUEST FROM REMOTE QUEUE  DEFAULT
      REORG     COMPRESS PDS DATA SETS                          DEFAULT
      RESETARC  DELETE ARCHIVE REQUEST FROM REMOTE QUEUE        DEFAULT
      RESETBKP  DELETE BACKUP REQUEST FROM REMOTE QUEUE        DEFAULT
      RESTAPPL  RESTORE DATA SETS FROM APPLICATION BACKUP      DEFAULT
      RESTARC   RESTORE DATA SETS FROM ARCHIVE                 DEFAULT
      RESTBKP   RESTORE DATA SETS FROM BACKUP                  DEFAULT
      SIMREORG  SIMULATE FDRREORG DATA SET REORGANIZATION      DEFAULT

```

This table defines the FDR and ABR commands that can be entered on the report output panels of SRS. On any line enter "S" to view or edit a command, "I" to insert a new command, "R" to copy and edit an existing command, or "D" to delete one. All except "D" take you to another panel, shown below, where you can view or modify the command.

54.51 CONTINUED . . .

Each FDR command consists of a main statement, a data set statement that is repeated for each selected data set, the program name to invoke and how to invoke it. The FDR Function commands are fully customizable. The user may change existing commands and define new commands with different keyword values in the main or data set statements. Keyword values may be automatically extracted from the Data Set List. The FDR Function commands may be saved in a private table library or in a common table library shared between users.

**FDR
FUNCTION
COMMAND
OPTIONS**

```

----- FDRSRS - FDR Function Command -----
COMMAND ==>
                                                                    SCROLL ==> PAGE

Command Name ==> COPY           Min number of characters for abbrev ==> 4   (0:8)
Description  ==> COPY DATA SETS

Main Control Statement:
==> COPY TYPE=DSF

Control Statement for Data Set Entry:   (omit DSN= and VOL= operands)
==> SELECT NOTIFY=&USERID

Default action after command is entered:
==> DISPLAY   ( DISPLAY | FG | RQ | SUBMIT | EDIT )

Request type ==> COPY           ( blank REORG COPY MOVE RESTART RESTBKP RESTAPPL
and/or                                     BACKUP ARCHIVE BACKAPPL other)
Program Name ==> FDRCOPY       (FDRABR FDRDSF FDRCOPY FDRABRUT)

Job Statement Information:   (only specify to override default)
==>
==>
==>
==>

SYSOUT Class ==>

TSO Userid group(s) of authorized users: (YES = all users; NO = no users)
FG ==> YES

BG ==> YES

RQ ==> NO

Recall dataset ==> YES         (yes no)   Suppress VOL= ==> NO         (yes no)
Setup CLIST    ==>                                     Cleanup CLIST ==>
Table Panel    ==>                                     JCL Skeleton   ==>
Pass entries   ==> NO         (yes no)   Menu Selection ==> DISPLAY (noaction)

Enter ADDISPF in the command line to make this command an ISPF command.
Enter ADDTSO  in the command line to make this command a TSO command.

Press PF3 (END) to save changes, or enter the CANCEL command to ignore changes.

```

This example shows the parameters for the COPY command. You can modify them if you like. Press PF3 (END) to save the changes.

Use the **ADDISPF** and the **ADDTSO** commands to add the FDR function command to the ISPF command table and as TSO command. This will enable the FDR function commands to be invoked as ISPF commands or as TSO commands in any panel in any application (including ISPF 3.4).

54.52 SRS DATA SET SELECTION CRITERIA PANEL

The selection, reporting and sorting criteria for data set selection and the source to be searched are specified in the Data Set Selection Criteria panel. On this panel, you specify the source for information about the data sets to be selected, give the tests that will be used to select those datasets, indicate how the data is to be sorted, and select the information fields that will be included in the display of the selected datasets.

PANEL A.S.1
SRS DATA SET
SELECTION

```

----- DATASET SELECTION: DEFAULT ----- LINE 1/181 COL 4:6/10
COMMAND ==>                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD          SELECTION VALUE                                REPORT SORT
-----
SOURCE FIELDS
DSNAME  ==>  -----
VOL      ==>  -----

SOURCE  ==>  CATALOG (Catalog Volume Archive Appl Scratch Extract)  ----

CATALOGN ==>  -----
ARCDN    ==>  -----
EXTDSN   ==>  FDRABR.SRS.EXTRACT -----

VTOC FIELDS
DEFAULTS ==>
UNIT      ==>  -----
UNITNAME  ==>  -----
DEVTYPE   ==>  -----
VOLSQ     ==>  -----
DATES      ==>  -----

```

The panel shown above is the default, which shows every field available to SRS. Since there are over 100 available fields, the panel is scrollable in 4 directions (UP, DOWN, LEFT, RIGHT). By scrolling UP (PF7/19) and DOWN (PF8/20) you can view all of the fields;. By scrolling LEFT (PF10/22) and RIGHT (PF11/23) you can view a command column on the left, and summary, summary break and field descriptions columns on the right (these are illustrated later).

Many users may never need to do selection on more than a data set name filter and/or volume serial mask. They will never need to use more than the first few fields and need not be concerned about scrolling or the other fields. But the power of SRS is available when they need it.

The **FIELD column** contains the name of a data set attribute (except for the SOURCE FIELDS, which specify the source to be searched). The fields available are essentially the same available for FDREPORT, the ABR Generalized Reporting Program, as described in [Section 54.31](#).

For those fields where the **SELECTION VALUE column** is underlined, you can specify a selection criteria based on that field (other fields are for reporting only and cannot be tested). To select on a given field, enter one or more values (separated by spaces) in the SELECTION VALUE column next to that field (the type of values depend on the nature of the field). Relational operators (= < > <= >= or EQ NE LT GT LE GE) can be inserted in front of each value (again separated by spaces) and supported in most fields. If a relational operator is not specified, = (equal) is the default.

For a data set to be selected by SRS, it must pass the tests for every Selection Value specified. If multiple Selection Values were specified for the same field, and no relational operators were specified (or they are all = or EQ), the test on that field will pass if it matches ANY of the values. If other relational operators are used, the field must match on ALL of the values.

54.52 CONTINUED . . .

The **REPORT column** indicates which report fields (data set attributes) should be displayed for the selected data sets, allowing you to totally customize the report. An S or a number (representing the order in which the fields are to be displayed) may be specified to select a field to be reported. If no REPORT fields are specified, then the Data Set List will contain the DEFAULTS fields (SPLDSN, VOL, DSORG, RECFM, BLKSIZE, LRECL, SIZE, SIZEFREE, %FREE).

The **SORT column** indicates which fields (data set attributes) should be sorted. This column does not apply when the Source is the ABR SCRATCH Catalog (in which case the Data Set List will be sorted by data set name). Not all fields can be selected for sorting. An S, a number (representing the order in which the fields are to be sorted) or a number followed by an A (for ascending) or a D (for descending) may be specified to select a field to be sorted. If more than one field is specified with an 'S', or if equal numbers are specified, those fields will be sorted in the order in which they appear in the Selection Criteria panel. If both 'S' and numbers are used in the Selection Criteria, then the fields with numbers will be sorted before the fields with 'S'. If no sorting is requested, the data sets will be in the order they were selected from the indicated source.

The **SUMMARY column** indicates which fields should be totaled for the selected data sets. An S or a number (representing the order in which the fields are to be displayed) may be specified to select a field to be summarized.

The **BREAK column** enables the user to select the fields that are to be sub-totaled and summarized in the Data Set List when the field value changes causing a control break. The break column is selected by specifying an "S" or the relative break position number.

A short list of Data Set Selection Criteria commands appears in the fourth line of the panel. These commands provide functions such as READING and SAVING Selection Criteria on disk, FINDING a string, LOCATING a field, SUBMITTING a batch job to process the Selection Criteria, etc. These commands may be entered on the COMMAND line at the top of the screen. However, simply pressing ENTER will execute the data set selection which has been specified on the panel, and will display the Data Set List panel with the data sets selected. A complete list and description of these commands is contained in the SRS HELP tutorials.

54.52 CONTINUED . . .

PANEL A.S.1
SCROLLED
LEFT

Scrolling LEFT (PF10/22) reveals the CMD column. Line commands such as D, DD (delete), R (repeat), M, MM (move), A (after) and B (before) may be entered to change the display panel.

These line commands can be used to customize which fields are displayed on the Selection panel and what order they are displayed in. This is used when you are developing a data set selection list to be saved for future use. Selection lists with a restricted set of commonly used fields might be saved in an installation-wide library for end-user use.

```

----- DATASET SELECTION: DEFAULT ----- LINE 1/181 COL 1:5/10
COMMAND ==>                                SCROLL ==> HALF

      Read   Save   Submit   Find   Locate   Extract   Options   Help

CMD  FIELD                SELECTION VALUE                                REPORT
----
---- SOURCE FIELDS
---- DSNAME   ==> -----
---- VOL      ==> -----
----
---- SOURCE   ==> CATALOG (Catalog Volume Archive Appl Scratch Extract) ----
----
---- CATALOGN ==> -----
---- ARCDN    ==> -----
---- EXTDSN   ==> FDRABR.SRS.EXTRACT_-----
----
---- VTOC FIELDS
---- DEFAULTS ==> -----
---- UNIT     ==> -----
---- UNITNAME ==> -----
---- DEVTYPE  ==> -----
---- VOLSQ    ==> -----
---- DATES    ==> -----

```

54.52 CONTINUED . . .

**PANEL A.S.1
SCROLLED
RIGHT**

Scrolling RIGHT (PF11/23) reveals the SUMMARY, BREAK, LEN and DESCRIPTION columns. LEN is the number of bytes that the field value will occupy in the Data Set List. DESCRIPTION is a brief description of the field. For some fields, such as SOURCE, it is necessary to press RIGHT several times to see all of the field DESCRIPTION. A detailed description of all fields is contained in the HELP tutorials.

----- DATASET SELECTION: DEFAULT --- LINE 1/181 COL 7:10/10									
COMMAND ==>					SCROLL ==> HALF				
	Read	Save	Submit	Find	Locate	Extract	Options	Help	
FIELD	SUMM BREAK LEN DESCRIPTION								
-----	----->								
SOURCE FIELDS									
DSNAME	==>	----	----	27	Data Set Name/VSAM cluster name				
VOL	==>	----	----	6	Volume Serial Data Set resides on				
SOURCE	==>	----	----	7	Source of input data: CATALOG (any or specified b				
CATALOGN	==>			44	Catalog Name where to search (optional)				
ARCDN	==>			44	Archive Control File data set name (optional)				
EXTDSN	==>			54	Data Set Name containing extracted data				
VTOC FIELDS									
DEFAULTS	==>			72	Includes: SPLDSN,VOL,DSORG,RECFM,BLKSIZE,LRECL,SI				
UNIT	==>	----	----	3	Device address the dataset is on.				
UNITNAME	==>			8	Esoteric or generic unit name				
DEVTYPE	==>	----	----	7	Type of Device Data Set resides on				
VOLSQ	==>	----	----	3	Data Set Volume Sequence Number				
DATES	==>			20	Includes: CRDATE,EXPDATE,LRDATE				

The default source to be searched is the CATALOG. Based on DSNAME, VOL and CATALOGN, data sets will be selected from the system catalogs, then the volsers indicated in the catalog will be accessed to extract additional fields relating to the data sets; if selection values for other fields were specified, the selected data sets will be filtered to eliminate those which do not match. The remaining datasets will be displayed with the requested (or default) report fields. If only DSNAME, VOL, and/or CATALOGN were selected on, you will have to press RIGHT (PF11/23) to view the other requested report fields.

For instance, to list all cataloged data sets starting with DF, either specify the high level index DF (which will automatically be converted into a data set filter), or the data set filter DF.** in the Selection Value column in the DSNAME row and press the ENTER key. The catalogs are searched and a Data Set List is displayed.

54.52 CONTINUED . . .

PANEL A.S.1
SIMPLE
CATALOG
EXAMPLE

----- DATASET SELECTION: DEFAULT ----- LINE 1/181 COL 4:6/10									
COMMAND ==>					SCROLL ==> HALF				
	Read	Save	Submit	Find	Locate	Extract	Options	Help	
FIELD	SELECTION VALUE						REPORT	SORT	

SOURCE FIELDS									
DSNAME	==>	DF.**						----	----
VOL	==>	-----						----	----
SOURCE	==>	CATALOG (Catalog Volume Archive Appl Scratch Extract)						----	
CATALOGN	==>	-----							
ARCDSN	==>	-----							
EXTDSN	==>	FDRABR.SRS.EXTRACT							

VTOC FIELDS									
DEFAULTS	==>	-----							
UNIT	==>	-----						----	----
UNITNAME	==>	-----							
DEVTYPE	==>	-----						----	----
VOLSQ	==>	-----						----	----
DATES	==>	-----						----	----

The data set name is automatically converted into a data set filter, unless it is imbedded in apostrophes, or already is a data set filter. The data set name filter uses the XDSN syntax of FDREPORT, documented earlier in this section. In simplest form, alphanumeric characters and periods represent themselves, * (one asterisk) represents any number of characters within a single index level, and ** (two asterisks) represent any number of characters in any number of index levels. Other special characters (such as + for any single numeric) are available.

Other sources are:

VOLUME	the data sets will be selected directly from the VTOCs of the indicated volumes; specify * on the VOL line to search all online volumes.
ARCHIVE	ARCHIVED data sets will be selected from an ARCHIVE Control File. By default, the common ACF indicated in the FDR/ABR Option table will be used, but any ACF can be specified by ARCDSN.
APPL	Application Backup data sets will be selected from the Control File. By default, the common ACF indicated in the FDR/ABR Option table will be used, but any Control file (See Section 52.08) can be specified by ARCDSN.
SCRATCH	selects from the ABR scratch catalog.
EXTRACT	reads an extract file produced by FDREPORT (see RPTYPE=DATA earlier in this section)

54.53 SRS DATA SET LIST PANEL

The SRS Data Set List Panel will be displayed when data sets are selected from a Data Set Selection panel, or when a saved Data set list is selected from the SRS primary panel. The format depends on the report fields selected on the original selection panel. If the selected report fields will not fit on one screen, you can scroll LEFT and RIGHT to view it all. If the selected data sets will not all fit on one screen, you can scroll UP and DOWN.

The following Data Set List Panel was displayed after a catalog search using the data set name filter DF.** as shown in the example in the previous section. The default report fields are shown.

**PANEL A.S.2
DATA SET LIST**

----- DATASET LIST: DSLIST ----- LINE 1/136 COL 3:9/10									
COMMAND ==>					SCROLL ==> HALF				
109 DATA SETS SELECTED.									
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help
COMMAND	DATA SET NAME				VOLSER	DSO	RECFM	BKSIZE	LRECL ALLOC FREE

	DF.\$IAM.CLUSTER				SMS812	*NO VOLUME DATA			
					*	*NO VOLUME DATA			
	DF.ABRMAP				MIGRAT ARCHIVED				
copy	DF.APF.LOAD				IDPLB1	PO	U	32760	146 52
copy	DF.ARCHIVE				IDPLB0	DA	F	23392	1 0
	DF.ASM.SOURCE				IDPLB1	PO	FB	19040	80 330 52
	DF.BROCHURE				MIGRAT ARCHIVED				
	DF.CALL.LOG				IDPLB1	PO	FB	6192	72 9 6
	DF.CAN				SCR083	*NO VOLUME DATA			
					*	*NO VOLUME DATA			
	DF.CMD.CLIST				IDPLB1	PO	FB	6240	80 60 18
	DF.CPK.REPORT				MIGRAT ARCHIVED				
	DF.CPK2.REPORT				MIGRAT ARCHIVED				
	DF.CSVLLA				IDPLB1	PO	FB	3360	80 1 0
	DF.DATA				IDPLB1	PO	FB	6240	80 15 0
	DF.DFP23.DOC				MIGRAT ARCHIVED				

Since the Data Set List column headers may be different from the selection field names (to occupy fewer columns). Column description may be obtained by positioning the cursor at the column and pressing the HELP key, or by entering the VIEW primary command.

Archived data sets are marked as ARCHIVED, and data sets not found on their cataloged volumes are marked as *NO VOLUME DATA. Unwanted data set entries in the Data Set List may be excluded from the display by entering the X or XX line commands in the COMMAND field. The Data Set List can be saved for redisplay at any time with option 2 on the SRS primary menu, and can be REFRESHed by regathering the report fields at any time.

Various services can be invoked by the user for one or more of the data sets displayed:

- SRS services (such as I (info), M (member list))
- FDR/ABR functions (such as Recall, Reorg, Copy)
- ISPF services (such as Edit, Browse)
- TSO commands (such as DELETE, LISTDS, LISTCAT, RENAME)
- CLISTs and REXX execs

by entering the command name in the COMMAND column next to the data set to be processed. Most commands need nothing more than the data set name, but some TSO commands and CLIST/REXX execs may need the data set name passed in a special way. If so, specify the complete command, with all operands, substituting a slash (/) where SRS is to substitute the data set name in apostrophes (You can type over the dsname on the panel; SRS will remember it). For example,

```
LISTCAT ENTRY(/) ALL
```


54.53 CONTINUED . . .

When you enter the name of a FDR/ABR service (such as COPY), operands valid for the service may also be included in the command. Action strings may also be included in the command, as follows:

- DISPLAY -- display a table for additional operands and execution
- FG -- execute the Service in the foreground (under TSO)
- RQ -- add request to the ABR remote queue immediately
- SUBMIT -- submit the generated JCL immediately
- EDIT -- edit the generated JCL

For instance, to copy a data set to another name in the foreground, specify

```
COPY / NEWINDEX=++XYZ FG
```

To perform the same function or service with the same operands on another data set further down in the Data Set List, enter the = (equal sign) repeat row command next to that data set.

Once a command is executed, it is displayed in the row COMMAND field preceded by an indicator representing the return code from the command: * (asterisk) for return code 0, ~ (not-sign) for return code 4, or ? (question mark) for all other return codes.

You can also specify a command to apply to all data sets displayed by entering it on the COMMAND line at the top of the screen. It must include a slash (/).

Data Set List commands, such as SAVE, READ, REFRESH, FIND, PRINTD, etc. may also be specified in the COMMAND line at the top of the screen. Please refer to the SRS HELP tutorial for a complete list and detailed description of the Data Set List commands.

The PRINTD (abbreviated P) command allows you to generate a printed report from the information in the Data Set List. A prompting panel will allow you to print all of the fields displayed, or to customize the report.

54.54 SRS DATA SET LIST FDR/ABR FUNCTIONS

The FDR/ABR functions supported on the Data Set List panel each have options that can be customized on the SRS option panels, as shown in [Section 54.51](#). When you enter a FDR/ABR function name, you may be prompted to override those defaults. For example, for a COPY function:

```
----- FDRSRS - Copy Data Set ----- ROW 1 TO 2 OF 2
COMMAND ==>                                SCROLL ==> HALF

      Edit generated JCL      Submit generated JCL      FG - execute in the foreground

Operands for COPY TYPE=DSF statement (section 21):
==> COPY TYPE=DSF

FROM DSNAME / Filter ==> 'DF.APF.LOAD'
      Volume Serial    ==> IDPLB1
TO   New DSNAME       ==> 'df.apf.load2'
      or NEWINDEX      ==>
New Volume Serial(s) ==>
Operands for SELECT DSN= statement (section 21):
==> NOTIFY=DF

FROM DSNAME / Filter ==> 'DF.ARCHIVE'
      Volume Serial    ==> IDPLB0
TO   New DSNAME       ==>
      or NEWINDEX      ==> ++copy2
New Volume Serial(s) ==>
Operands for SELECT DSN= statement (section 21):
==> NOTIFY=DF
```

The defaults for each FDR/ABR function can be modified using **option 0** on the SRS Primary Menu, or the FUNCTION (FF or FS) command can be used on the COMMAND line of the Data Set List panel. For convenience, you can go directly to the options panel for a particular function, for instance: FF COPY

54.55 SRS VOLUME SELECTION CRITERIA PANEL

The selection, reporting and sorting criteria for volume selection are specified in the Volume Selection Criteria panel. On this panel, you give the tests that will be used to select volumes, indicate how the volume data is to be sorted, and select the information fields that will be included in the display of the selected volumes.

PANEL A.S.3
VOLUME
SELECTION
CRITERIA

```

----- VOLUME SELECTION: DEFAULT ----- LINE 1/72 COL 4:6/10
COMMAND ==>                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD      SELECTION VALUE                                REPORT SORT
-----
VLVOLSER ==> PROD*-----
VLUNIT  ==> -----
UNITNAME ==> -----
VLDEVTYP ==> -----
STORGRP  ==> -----

VL%FTRKS ==> -----
VL%UTRKS ==> -----
VL%FDSCB ==> -----
VL%UDSCB ==> -----
VL%FINDX ==> -----
VL%UINDX ==> -----
VL%FVVDs ==> -----
VL%UVVDs ==> -----
VLALOTRK ==> -----
VLAVGTRK ==> -----
VLALOSTA ==> -----

```

The Volume Selection Criteria Panel functions in a similar fashion to the Data Set Selection Criteria Panel documented in the preceding sections. There are more fields which can be viewed by scrolling UP and DOWN. The command field and field descriptions can be viewed by scrolling LEFT and RIGHT. Selection, reporting, summary and sorting criteria are specified just like the Data Set Selection Criteria.

Volume selection differs from data set selection in that one line of the generated volume list contains fields relating to one whole DASD volume. There is no SOURCE field since there is no choice of source with volume selection; information is gathered from the VTOC, VTOCIX, and VVDS of the volumes selected, and the LSPACE SVC.

54.56 SRS VOLUME LIST PANEL

The following Volume List was generated by specifying PROD* in the VLVOLSER field in the Volume Selection Criteria. The volume list shown is the default if no report fields are selected.

PANEL A.S.4
VOLUME LIST

----- VOLUME LIST: VOLLIST ----- LINE 1/13 COL 3:12/14											
COMMAND ==>						SCROLL ==> HALF					
13 VOLUMES SELECTED.											
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help		
COMMAND	VOLSER	DEVTYPE	UADR	VTINDX	USEATTR	USERS	ALOTR	%TU	FRTRK	LFCYL	LFTRK

	PRODM1	3380	0140	ACTIVE	PRIVATE	0	11979	91	1296	41	640
	PRODB0	3380	0141	ACTIVE	PRIVATE	19	11083	84	2192	144	2160
	PRODK1	3380	0146	ACTIVE	PUBLIC	0	2612	20	10663	684	10273
	PRODB3	3380	0147	ACTIVE	PRIVATE	62	12235	93	1040	34	526
	PRODB2	3380	0149	ACTIVE	PRIVATE	42	11934	90	1341	78	1170
	PRODB1	3380-K	0151	ACTIVE	STORAGE	205	37238	94	2587	58	884
	PRODB5	3380-K	0152	YES	STORAGE	309	33202	84	6623	198	2970
	PRODB4	3380-K	0154	ACTIVE	STORAGE	195	35433	89	4392	287	4310
	PRODM3	3380-K	0156	ACTIVE	PRIVATE	13	27678	70	12147	807	12118
	PRODM0	3380	0183	ACTIVE	PRIVATE	198	13129	99	146	9	142
	PRODK0	3380	0184	ACTIVE	STORAGE	1	6951	53	6324	317	4758
	PRODK2	3380	0186	ACTIVE	STORAGE	3	7591	58	5684	190	2864
	PRODM2	3380	0188	ACTIVE	PRIVATE	76	13118	99	157	1	15

Since the Volume List column headers may be different from the selection field names (to occupy fewer columns), column descriptions may be obtained by positioning the cursor at the column and pressing the HELP key, or by entering the VIEW primary command.

Since the volume column headings may not be immediately obvious, here is what some of them mean (the last 2 are not shown on the panel above since you must scroll right to see them:

ALOTR -- total allocated tracks on the volume

%TU -- percentage of tracks allocated

FRTRK -- total free tracks on volume

LFCYL -- largest free cylinder extent on the volume

LFTRK -- largest free track extent on the volume

FREXT -- number of free extents on the volume

FRAG -- IBM fragmentation index for the volume

You may enter an **S** in the command column to display the SRS data set list for all the data sets on the selected volume.

54.57 SRS ISPF FASTPATHS AND COMMANDS**ISPF
FASTPATHS**

Once you have become familiar with using SRS, you can use these ISPF "fastpath" options from the ISPF main menu to invoke SRS without going through the intermediate panels:

A.S. -- executes the DEFAULT Data Set Selection Criteria and selects data sets matching your TSO prefix or TSO userid.

A.S.dsfilter -- executes the DEFAULT Data Set Selection Criteria, and selects the data sets matching the data set name filter specified. The filter can be any type of data set filter supported by the XDSN= operand of FDREPORT as documented earlier in this section. If the data set filter is not in quotes, SRS will assume that it is a prefix and append ".*" to the end. For fully-qualified names, place them in single quotes. For example:

```
A.S.TSO1
A.S.'TSO1.JCL.CNTL'
A.S.'TSO+.*.CNTL'
```

A.S.selname,dsfilter -- same as the previous fastpath except that the Data Set Selection Criteria saved under name "selname" will be used. For example:

```
A.S.ABRBKUP,TSO1
```

In all of the preceding A.S fastpaths, you may optionally follow them with a backslash and a volume serial or volume prefix to limit the display to those volumes. For example:

```
A.S.TSO1\PUB*
A.S.ABRBKUP,PROD.PAYROLL\PAY123
A.S.\WORK*
```

A.S1 -- displays the most recently used saved Data Set Selection Criteria, allowing you to modify and execute it.

A.S1.selname -- displays the named saved Data Set Selection Criteria. For example,
A.S1.ABRBKUP

A.S2 -- displays the most recently used saved Data Set List.

A.S2.listname -- displays the named Data Set List.

A.S3 -- displays the most recently used Volume Selection Criteria list.

A.S3.volser -- executes the most recently used Volume Selection Criteria list against the volume serial or volser prefix specified. For example,

```
A.S3.TSO123
A.S3.TSO1*
```

A.S3.selname, -- displays the named Volume Selection Criteria list.

A.S3.selname,volser -- executes the named Volume Selection Criteria list against the volume serial or volser prefix specified.

A.S4 -- displays the most recently used saved Volume List.

A.S4.listname -- displays the named saved Volume List.

Since ISPF will also allow you to stack ISPF commands on one line, separated by a semicolon (by default), when the ISPF fastpath results in a data set or volume list being displayed, you can also specify a command to be executed against all of the displayed data sets or volumes. For example:

A.S.PDS,USER1;REORG / FG -- assuming that PDS is a saved selection criteria that selects only PDSs, this will execute a reorganization (PDS compression) against every PDS belonging to USER1.

**ISPF
COMMANDS**

You can invoke SRS or an FDR function as a primary ISPF command from any panel in any ISPF application. [See Section 90.46](#) "Installing the FDR Dialogs" for details.

54.60 SRS EXAMPLES

This section "walks" you through several examples showing the ease of use and power of SRS for several common functions. The display at each step in each example is shown; data in italics was entered on that panel by the user.

**RESTORE
FROM
ARCHIVE**

A user wishes to display all of his ARCHIVED data sets and select several for restore. From the SRS main menu, option 1 (select data sets) with the default selection criteria was chosen:

```

----- FDRSRS - Primary Menu -----
OPTION  ===> 1                                SCROLL ===> HALF
                                           More:      +

  0  OPTIONS  - Set Dialog Options and Defaults

  1  SELECT   - Data Set Selection
              Name ===> DEFAULT              (*, member name, or blanks)

  2  DSLIST   - Data Set List - display saved
              Name ===> *                    (*, member name, or blanks)

  3  SELVOL   - Volume Selection
              Name ===> DEFAULT              (*, member name, or blanks)

  4  VOLLIST  - Volume List - display saved
              Name ===> *                    (*, member name, or blanks)

OR Select one of the following services:
_  ARCDL     - DELETE DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCHIVE   - ARCHIVE DATA SET
_  ARCMOD    - MODIFY DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCRCAT   - RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL
_  ARCRES    - RESET DATA SET ENTRY IN THE ARCHIVE FILE
_  BACKAPPL  - DATA SET APPLICATION BACKUP
----- Copyright 1992, 1998 --- Innovation Data Processing, Inc. -----

```

54.60 CONTINUED . . .

On the selection panel, a data set prefix is entered, and the source is changed to ARCHIVE. If you need to report on a special ARCHIVE Control File or a Control File used for Application Backup, you can enter the control file name under ARCDN; if left blank, the ARCHIVE control file used for auto-recall (from the FDR Global Option Table) is automatically used. Since no reporting criteria are specified, the default report is displayed.

```

----- DATASET SELECTION: DEFAULT ----- LINE 1/181 COL 4:6/10
COMMAND ==>                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD          SELECTION VALUE                                REPORT SORT
-----
SOURCE FIELDS
DSNAME  ==>  grp.**-----
VOL      ==> -----
SOURCE  ==>  archive (Catalog Volume Archive Appl Scratch Extract) -----
CATALOGN ==> -----
ARCDN    ==> -----
EXTDSN   ==> -----

VTOC FIELDS
DEFAULTS ==> -----
UNIT      ==> -----
UNITNAME  ==> -----
DEVTYPE   ==> -----
VOLSEQ    ==> -----
DATES      ==> -----
CRDATE    ==> -----
CRDAYS    ==> -----

```

A list of ARCHIVED data sets that match the data set filter are displayed with their characteristics. The list can be scrolled up and down to see all the data sets, and left and right to see all of the fields that SRS displays by default. If you want to restore any of them, simply type "restore" next to them.

```

----- DATASET LIST: DSLIST ----- LINE 73/420 COL 3:9/13
COMMAND ==>                                SCROLL ==> HALF
420 DATA SETS SELECTED.
      Read   Save   Find   Locate   Refresh   Next   Message   Printd   Help

COMMAND  ENTRY NAME                                VOLSER DSO RECFM BKSIZ LRECL ALLOC  FREE
-----
restore  QRP.FLOR.INFO                                IDPLB3 PS  FB    3120   80    2    0
restore  QRP.DEFRAG.TXT                                IDPLB4 PS  FB    9040   80    1    0
         QRP.IAM.THIRD                                IDPLB3 PS  FB    3120   80    5    0
         QRP.V2R3.TABLES                                IDPLB4 PO  FB    6160   80    1    0
restore  QRP.JCL.CNTL                                IDPLB3 PO  FB    3120   80   90   19
         QRP.DASD.D33909                                IDPLB1 PS  FBA   8000   80    1    0

```

54.60 CONTINUED . . .

The "restore" command displays a panel where options for the restore, such as new name or new volser, can be specified for each of the selected data sets. If more data sets were selected than can fit on the screen, the list is scrollable. Enter a command on the command list to submit the restore as a batch job (or edit the JCL before submission), execute the restore immediately under TSO, or add it to the ABR remote queue for later processing.

```
----- FDRSRS - Archive Restore ----- Row 1 to 2 of 2
COMMAND ==> submit                                SCROLL ==> HALF

  Edit JCL   Submit JCL   FG - execute in the foreground   RQ - add to remote q

Operands for RESTORE TYPE=ARC statement (section 51.06):
==> RESTORE TYPE=ARC,DT,DYNARC

DSNAME / Filter ==> 'QRP.DEFRAG.TXT'
Volume Serial   ==> IDPLB4                               Archive date ==>
New DSNAME      ==>
or NEWINDEX     ==> .defrag2
New Volser(s)   ==>                                       Copy ==>
Operands for SELECT DSN= statement (section 51):
==> NOTIFY=BAB

DSNAME / Filter ==> 'QRP.JCL.CNTL'
Volume Serial   ==> IDPLB3                               Archive date ==>
New DSNAME      ==>
or NEWINDEX     ==>
New Volser(s)   ==> tsowk*                               Copy ==>
Operands for SELECT DSN= statement (section 51):
==> NOTIFY=BAB
```


54.60 CONTINUED . . .

**RESTORE
FROM BACKUP**

A user needs to restore several of his data sets from ABR backups. From the SRS main menu, option 1 (select data sets) was chosen and a selection name of "*" is specified to list all available saved selection lists:

```

----- FDRSRS - Primary Menu -----
OPTION  ==> 1                                SCROLL ==> HALF
                                           More:      +

  0  OPTIONS  - Set Dialog Options and Defaults

  1  SELECT   - Data Set Selection
                Name ==> DEFAULT              (*, member name, or blanks)

  2  DSLIST   - Data Set List - display saved
                Name ==> *                    (*, member name, or blanks)

  3  SELVOL   - Volume Selection
                Name ==> DEFAULT              (*, member name, or blanks)

  4  VOLLIST  - Volume List - display saved
                Name ==> *                    (*, member name, or blanks)

OR Select one of the following services:
_  ARCDL     - DELETE DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCHIVE   - ARCHIVE DATA SET
_  ARCMOD    - MODIFY DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCRECAT  - RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL
_  ARCRESET  - RESET DATA SET ENTRY IN THE ARCHIVE FILE
_  BACKAPPL  - DATA SET APPLICATION BACKUP
----- Copyright 1992, 1998 --- Innovation Data Processing, Inc. -----

```

The list of saved selection and reporting criteria is displayed. Depending on the SRS options set, these saved criteria might have been previously created by the user, or might be criteria available to all SRS users, or both. The ABRBKUP list is selected, since it will display all information about the ABR backups of selected data sets:

```

----- FDRSRS - Selection Criteria List - Row 1 to 15 of 15
Command ==>                                Scroll ==> HALF

Place an "S" to select the Selection Criteria to process, or "D" to delete.

Sel Member  Description                                Date
-----
s  ABRBKUP   abr backup info                            06/11/1996
   ARCBKUP   archive backup info                        06/11/1996
   ARCEF     archived icf datasets                      10/20/1992
   BACKEXT   report from backup extract                 09/10/1996

```

54.60 CONTINUED . . .

The selection panel is now displayed, preset with the selection and reporting criteria from the ABRBKUP list. In this case, ABRBKUP simply defines a report format, so only REPORT fields are pre-specified; the numbers shown in the REPORT column is the order that those fields will be displayed on the screen. The list has been customized so that only the selected report fields are displayed. Several data set name filters are entered to display backup info about those data sets.

```

----- DATASET SELECTION: ABRBKUP ----- LINE 1/16 COL 4:6/10
COMMAND ==>                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD          SELECTION VALUE                                REPORT SORT
-----
DSNAME ==> bab.*.jcl bab.ac**                                1___ ____
VOL ==> _____                                2___ ____

SOURCE ==> CATALOG (Catalog Volume Archive Appl Scratch Extract)  ____ ____

CATALOGN ==> _____

DEVTYPE ==> _____                                3___ ____
DSORG ==> _____                                4___ ____
BKGEN ==> _____                                6___ ____
BKCYCLE ==> _____                                7___ ____
BKINFO ==> _____                                8___ ____
OLDBKUP ==> ALL _____                                ____ ____
SIZE ==> _____                                5___ ____

```

A list of the selected data sets is displayed, along with backup information for each. For those that have more than one backup recorded (the OLDBACKUP option), all of the recorded backups are displayed, including the date that each backup was taken. Since this selected data sets from the system catalogs, archived data sets may also be selected, as shown. To restore a data set from backup, simply type "restore" on the line for the backup desired:

```

----- DATASET LIST: DSLIST ----- LINE 1/20 COL 3:9/12
COMMAND ==>                                SCROLL ==> HALF
3 DATA SETS SELECTED.
      Read   Save   Find   Locate   Refresh   Next   Message   Printd   Help

COMMAND  DATA SET NAME                                VOLSER DEVTYPE DSO  ALLOC BKGn BKC    BKDATE
-----
restore  BAB.JCL.CNTL                                TSOWK2 3390-2  PO    150    3    5 03/13/1998
                                                3    4 03/12/1998
restore  BAB.AC.DATA                                IDPLB3 3390-2  PS     1  764    0 03/06/1998
                                                763  18 03/05/1998

```

54.60 CONTINUED . . .

The "restore" command displays a panel where options for the restore, such as new name or new volser, can be specified for each of the selected data sets. Note that the correct volume, gen and cycle to restore the backup selected are already filled in. If more data sets were selected than can fit on the screen, the list is scrollable. Enter a command on the command list to submit the restore as a batch job (or edit the JCL before submission), execute the restore immediately under TSO, or add it to the ABR remote queue for later processing.

```

----- FDRSRS - Backup Restore ----- Row 1 to 2 of 2
COMMAND ==> edit                                SCROLL ==> HALF

  Edit JCL    Submit JCL    FG - execute in the foreground    RQ - add to remote q

Operands for RESTORE TYPE=ABR statement (section 50):
==> RESTORE TYPE=ABR,DT

DSNAME/Filter ==> 'BAB.JCL.CNTL'                    Gen ==> 0003
Volume Serial ==> TSOWK0                            Cycle => 05
New DSNAME     ==>                                OLDB =>
or NEWINDEX    ==>                                Copy  => 1
New Volser(s) ==>
Operands for SELECT DSN= statement (section 50):
==> NOTIFY=BAB

DSNAME/Filter ==> 'BAB.AC.DATA'                    Gen ==> 0764
Volume Serial ==> IDPLB3                            Cycle => 00
New DSNAME     ==>                                OLDB =>
or NEWINDEX    ==>                                Copy  => 1
New Volser(s) ==>
Operands for SELECT DSN= statement (section 50):
==> NOTIFY=BAB

```

54.60 CONTINUED . . .

VOLUME REPORTS A Storage Analysis wants to monitor usage and free space on production and data base volumes. From the SRS main menu, option 3 (select volumes) with the default selection criteria was chosen:

```

----- FDRSRS - Primary Menu -----
OPTION  ==> 3                                SCROLL ==> HALF
                                           More:      +

  0  OPTIONS  - Set Dialog Options and Defaults

  1  SELECT   - Data Set Selection
                Name ==> DEFAULT              (*, member name, or blanks)

  2  DSLIST   - Data Set List - display saved
                Name ==> *                    (*, member name, or blanks)

  3  SELVOL   - Volume Selection
                Name ==> DEFAULT              (*, member name, or blanks)

  4  VOLLIST  - Volume List - display saved
                Name ==> *                    (*, member name, or blanks)

OR Select one of the following services:
_  ARCDL     - DELETE DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCHIVE   - ARCHIVE DATA SET
_  ARCMOD    - MODIFY DATA SET ENTRY IN THE ARCHIVE FILE
_  ARCRECAT  - RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL
_  ARCRESET  - RESET DATA SET ENTRY IN THE ARCHIVE FILE
_  BACKAPPL  - DATA SET APPLICATION BACKUP
----- Copyright 1992, 1998 --- Innovation Data Processing, Inc. -----

```

On the selection panel, volume serial prefixes are entered. Since no reporting criteria are specified, the default report is displayed.

```

----- VOLUME SELECTION: DEFAULT ----- LINE 1/72 COL 4:6/10
COMMAND ==>                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD      SELECTION VALUE                                REPORT SORT
-----
VLVOLSER ==> PROD* DBLG*                                -----
VLUNIT  ==> -----                                -----
UNITNAME ==> -----                                -----
VLDEVTYP ==> -----                                -----
STORGRP  ==> -----                                -----

VL%FTRKS ==> -----                                -----
VL%UTRKS ==> -----                                -----
VL%FDSCB ==> -----                                -----
VL%UDSCB ==> -----                                -----
VL%FINDX ==> -----                                -----
VL%UINDX ==> -----                                -----
VL%FVVDX ==> -----                                -----
VL%UVVDX ==> -----                                -----
VLALOTRK ==> -----                                -----

```

54.60 CONTINUED . . .

The status of the selected volumes is displayed, including Indexed VTOC status, mount status, tracks allocated, percentage used, and free space information. If you scroll right, the number of free extents and the IBM fragmentation index for each volume is visible.

----- VOLUME LIST: VOLLIST ----- LINE 1/7 COL 3:12/14											
COMMAND ==>						SCROLL ==> HALF					
7 VOLUMES SELECTED.											
	Read	Save	Find	Locate	Refresh	Next	Message	Printd	Help		
COMMAND	VOLSER	DEVTYPE	UADR	VTINDX	USEATTR	USERS	ALOTR	%TU	FRTRK	LFCYL	LFTRK

	DBLG01	3380	0143	ACTIVE	STORAGE	115	2934	23	10341	438	6573
	PROD01	3380-K	0151	ACTIVE	STORAGE	238	34027	86	5798	187	2811
	PROD05	3380-K	0152	ACTIVE	STORAGE	448	38726	98	1099	28	420
	PROD04	3380-K	0154	ACTIVE	STORAGE	251	37238	94	2587	82	1230
	PROD03	3380-K	0155	ACTIVE	PRIVATE	83	35703	90	4122	180	2700
	DBLG00	3390-3	0180	ACTIVE	STORAGE	129	46340	93	3745	53	801
	PROD00	3390-3	0183	ACTIVE	PRIVATE	12	49705	99	380	24	360

54.60 CONTINUED . . .

**FASTPATH
VOLUME
DISPLAY**

Operations wants to regularly check on the free space available on various volumes, to anticipate and avoid allocation problems. Previously, a SRS volume report was customized and saved as FREESPC, as shown here. It displays the percentage of free space on the volume and in the VTOC as well as total free tracks on each volume. It only selects volumes which have less than 20% free space.

```

----- VOLUME SELECTION: DEFAULT ----- LINE 1/7 COL 4:6/10
COMMAND ==> save freespc                                SCROLL ==> HALF
ENTER SELECTION CRITERIA
      Read   Save   Submit   Find   Locate   Extract   Options   Help

FIELD          SELECTION VALUE                                REPORT SORT
-----
VLVOLSER ==> ----- S____
VLUNIT  ==> ----- S____
VLDEVTYP ==> ----- S____

VL%FTRKS ==> <20----- S____
VL%FDSCB ==> ----- S____
VLFRETRK ==> ----- S____

```

Once this is saved, the operators can request the report from the ISPF main menu, specifying the volume serial or volser prefix to be displayed:

```

Menu  Utilities  Compilers  Options  Status  Help
-----
                                ISPF Primary Option Menu
Option ==> a.s3.freespc,sys*

0  Settings      Terminal and user parameters      <   Calendar   >
1  View          Display source data or listings      March      1998
2  Edit          Create or change source data      Su Mo Tu We Th Fr Sa
3  Utilities     Perform utility functions          1  2  3  4  5  6  7
4  Foreground    Interactive language processing      8  9 10 11 12 13 14
5  Batch         Submit job for language processing    15 16 17 18 19 20 21
6  Command       Enter TSO or Workstation commands      22 23 24 25 26 27 28
7  Dialog Test   Perform dialog testing          29 30 31
8  LM Facility   Library administrator functions
9  IBM Products  IBM program development products      Time . . . . : 12:10
10 SCLM          SW Configuration Library Manager      Day of year. :   082
A  FDR/ABR       FDR/ABR DASD Management Functions

```

and receive a report such as:

```

----- VOLUME LIST: VOLLIST ----- LINE 1/3 COL 3:7/7
COMMAND ==>                                SCROLL ==> HALF
3 VOLUMES SELECTED.
      Read   Save   Find   Locate   Refresh   Next   Message   Printd   Help

COMMAND  VOLSER  UADR  DEVTYP  %TF  %DF  FRTRK
-----
          SYSLB4 016B 3390-2   19  36  6424
          SYLOS1 016C 3390-2   16  37  5469
          SYSLB3 016E 3390-2    9  63  3116

```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

60.01 FDR AND ABR BACKUP MAINTENANCE

Once backups on tape or disk have been created with FDR, FDRDSF, or FDRABR, you may need to do maintenance on those backups. This maintenance may be as simple as creating additional copies of FDR backups, or complex operations such as deleting obsolete ABR ARCHIVE backups.

Two utility programs are provided for backup maintenance:

- FDRTCOPY can be used to copy any FDR or ABR backup, and also performs special maintenance functions on ABR backups. It is included even for sites licensed only for FDR.
- FDRTSEL is used only with ABR; it automates copying and condensing of ABR backups.

WARNING: Backups created by FDR, FDRDSF, SAR, and FDRABR are in a unique FDR format and cannot be correctly copied by programs other than FDRTCOPY and FDRTSEL. If you copy FDR-format backups with any other copy utility, such as the IBM utility IEBGENER, the copy may appear to be successful but any attempt to restore from the copied backup will fail. Similarly, programs which transmit files from one computer site to another will usually corrupt FDR backups. The FATAR Tape utility (another Innovation product) can also copy FDR backup tapes. If you must copy FDR tapes with other copy or transmission utilities, please contact Innovation for assistance.

**COPYING
FDR BACKUPS**

During an actual backup, you can create 1 or 2 copies of the backup file simultaneously. However, it may be more convenient to create the second copy at a later time, or you may want more than 2 copies. FDRTCOPY can do this.

All backups created by FDR, FDRDSF, FDRABR, or SAR have the same internal format. In its simplest usage, the FDRTCOPY utility can be used to make an exact copy of any FDR-format backup. It can also be used to copy backups on disk to tape, or vice versa.

While doing the copy, FDRTCOPY will also validate the internal format of the backup, verifying the structure of the data blocks and also verifying that all data which is supposed to be on the tape is actually present and readable.

The data set names and locations of non-ABR backups are totally under the user's control. The FDRTCOPY JCL for copying non-ABR tapes must point to the input backup, and the name and location of the copy is also specified in the JCL.

**COPYING
ABR BACKUPS**

Backups created by FDRABR use special naming conventions, and those backups are recorded by ABR, in the ABR catalog for full-volume and incremental backups, and in the Archive Control File for ARCHIVE backups and application backups. The FDRTCOPY and FDRTSEL utilities have special support for copying ABR backups. They recognize the ABR backup data set names, can copy multiple backup files from an input backup, and will update the appropriate ABR records to record the copies.

FDRTSEL automates the copy process, looking up the selected backups in the ABR catalog or Archive Control File and dynamically allocating the input backups.

FDRTCOPY can also be used to directly copy ABR backups, but the user must point, via JCL, to the backup file to be copied (or the first file on a backup tape). Innovation recommends the use of FDRTSEL for most ABR backup copying.

60.01 CONTINUED

**ABR ARCHIVE
MAINTENANCE**

FDRTSEL has other important functions for ABR ARCHIVE backups.

- It automates the movement of expired (or about to expire) ARCHIVE backups on disk to tape. This is a key step if you have structured your ARCHIVE jobs to place the first copy on disk with a short term retention (for quick auto-recall) but want to move that copy to tape with a longer retention when it reaches its first expiration.
- It can copy your existing ARCHIVE backup tapes and eliminate the backups of disk data sets which have expired and are no longer needed. The output copy will be significantly smaller than the original since it will contain only the remaining active disk data sets, and will free up tapes in your library.
- It may be used for migration from one type of tape to another, such as 3480 to 3490E or Magstar (3590). Since the newer tapes hold much more data on a volume, this can also greatly reduce the number of tape volumes required to hold your archival data.

60.10 FDR TAPE COPY (FDRTCOPY)

OVERVIEW The FDR tape copy utility (FDRTCOPY) has been specifically designed to copy FDR-formatted backups on tape or disk. The FDR format is used by all backups created by FDR, FDRDSF, FDRABR, or SAR (Stand-Alone Restore). **As noted in [Section 60.01](#), FDR backups cannot be correctly copied by any non-Innovation utilities.**

FDRTCOPY may be used to copy any FDR backup, from disk or tape, to disk or tape, optionally creating a second copy at the same time. Backups in FDR compressed format can be copied, and FDRTCOPY can create compressed output from uncompressed input, and vice versa. Simple JCL and control statements are used to specify the input and output backup data sets.

**COPYING
ABR BACKUPS**

When the input is ABR-created backup or archive files, FDRTCOPY has the ability to:

- Automatically copy multi-file tapes (tapes containing multiple backups or archives stacked on tape by ABR).
- Create replacement copies of backups, i.e., create and catalog data sets with the same name as the input.
- Create duplicate copies of backups, i.e., create COPY2 backups from COPY1 input, or vice versa.
- Create additional copies (COPY3 through COPY9) of full-volume and incremental backups.
- Drop expired backup files during a multi-file copy, allowing reduction and consolidation of incremental and archive files.
- Update the ABR Archive Control File if archive files are copied.

FDRTCOPY will automatically recognize that input backups were created by ABR based on their format of their data set names. The dsname format of ABR-created backups is described in detail in [Section 52.05](#), but briefly it is:

abrindex.Vvolser.xczzzzzz

where "abrindex" is a fixed index (usually "FDRABR"), "volser" is the volume serial of a disk volume, "x" is "C" for backups and other characters for archives and application backups, "c" is a copy number (1-9), and "zzzzzz" is variable data. FDRTCOPY gets the usual "abrindex" from the FDR/ABR Global Option Table, but note that the "abrindex" may be any user-chosen value for Application Backups (DUMP TYPE=APPL).

FDRTCOPY will do a pattern-match on the input data set name to see if it meets the requirements for an ABR backup (if you do not wish FDRTCOPY to identify ABR backups, e.g., if you want to copy an ABR backup to a non-ABR backup name, specify "ABR=NO" on the COPY control statement). When the input is on tape, the tape label only contains the last 17 characters of the dsname, so FDRTCOPY cannot verify the "abrindex" and checks only the remainder. If an input file is identified as ABR-created, the following rules apply ([see Section 60.12](#) for the meaning of the FDRTCOPY operands):

- the input data set name will always be copied to the output data set (COPYDSN=YES is assumed). The output name will be identical to the input unless ABRCOPY= and/or ABRCOPY2= is specified; these operands allow you to change the copy number "c" in the dsname.
- the output data sets will be cataloged unless CAT= and/or CAT2= is specified (CAT=RECAT is the default).
- if the input is identified as an archive backup, FDRTCOPY will attempt to update the Archive Control File to point to the output backup.
- if the end of the last input tape is reached without finding the end of the backup data (as indicated by internal FDR control blocks), FDRTCOPY will attempt to locate the backup in the ABR catalog to get the remainder of the backup tape volume serials. This allows you to copy an ABR tape set while specifying only the first input tape (or first few tapes) in the JCL.
- If ALLFILES or MAXFILES= is also specified, FDRTCOPY will copy multiple ABR files from the input tape(s). Since ABR-created backups are normally multi-file, this allows all the backups on a tape set to be copied without specifying anything except the first file to be copied. If necessary, it will use the ABR catalog to get the additional tape volsers for each backup, until either the number of files specified by MAXFILES= has been reached, or a double tape mark (indicating the end of the tape set) is encountered.

Files which are not identified as ABR-created will not receive this special processing.

60.10 CONTINUED

**COPYING
BACKUP
SUBSYSTEM
TAPES**

FDRTCOPY can be used to copy ABR full-volume and incremental backup tapes. The JCL must point to the input backup (or the first backup on a multi-file tape). Given the nature of the backups created by ABR, it is not always easy to make up standard JCL to copy the tapes created by ABR every day, but [Section 60.13](#) contains some sample JCL that can be used to do so.

This might be used to create a duplicate copy in a separate run, rather than using the TAPExx facility of ABR to create the duplicate at dump time; you might do this if you don't have enough tape drives or if you are trying to reduce the dump elapsed time. It can also be used to create additional duplicate copies for offsite storage; ABR supports restoring from copy numbers 3 through 9, but they can be created only by FDRTCOPY. FDRTCOPY can also be used to recreate a backup from a duplicate copy when the original has been damaged or lost.

When copying backup files, FDRTCOPY always copies the ABR-created data set name (See [Section 52.05](#)), but the copy number in the name can be changed by using the ABRCOPY=/ABRCOPY2= operands.

Since ABR uses only the ABR catalog to locate backup subsystem tapes, it is essential that the output files created by FDRTCOPY be cataloged. The default of CAT=RECAT and CAT2=RECAT will ensure that the output files are always cataloged; this is especially important when you are recreating tapes which are already cataloged to different volumes.

FDRTSEL can also be used to automate the copying of ABR incremental and full volume backups. FDRTSEL can select the backups from the ABR catalog and dynamically allocate the inputs, eliminating the FDRTCOPY requirement for JCL pointing to the backups. FDRTSEL also has options for ordering the backups to improve restore performance. See [Section 60.20](#).

**COPYING
ARCHIVE
SUBSYSTEM
TAPES**

FDRTCOPY can be used to copy ABR archive backup tapes. The JCL must point to the input backup (or the first backup on a multi-file tape); [Section 60.13](#) contains some sample JCL that can be used to automate these copies.

Archive backups are recorded in the Archive Control File. They may also be cataloged in the ABR catalog, but ABR does not use the catalog entries for restore unless the backup exceeds 5 tape volumes. The structure of the Archive Control File only allows copy numbers 1 and 2 to be recorded; when FDRTCOPY copies an archive backup file and the output file copy number is 1 or 2, it updates the Archive Control File using this technique:

- The ID of the input file (data set name, tape volser, and file sequence number) are noted.
- The ID of the equivalent output file (name, tape volser, file sequence number, expiration date and device type) are also noted.
- The Archive Control File is searched for any archive data set which is identified as being archived on the input file (both COPY1 and COPY2 are checked against the input file ID).
- When a match is found, the ID of the output file is inserted into the slot for COPY1 or COPY2, depending on which copy number appears in the output file name.

With this technique, you can create a COPY2 from a COPY1 input, and FDRTCOPY will find the proper entries from the COPY1 name and add the COPY2 ID. If you are copying an archive tape using the same copy number, FDRTCOPY will replace the old backup IDs with the new volume serials, device type, and expiration.

FDRTCOPY will only update the Archive Control File if you point to that file, by providing an ARCHIVE DD statement, or, if the file to be updated is the one named in the ABR global option table, by specifying the DYNARC operand.

60.10 CONTINUED

**MOVING
ARCHIVE FILES
FROM DISK**

Although FDRTCOPY is capable of moving archive files from disk to tape, the procedure is laborious so it is not convenient to do so. FDRTSEL is a utility which can automate the process of copying disk archive files to tape, invoking FDRTCOPY to do the actual move. FDRTSEL is especially useful if you archive to disk with short-term retention, and want to create tape copies of those files when they expire.

FDRTCOPY is also capable of moving archive backups from disk to disk, but again the process is laborious. FDRTSEL can also automate this process.

Documentation on FDRTSEL is found in [Section 60.20](#).

**APPLICATION
BACKUP FILES**

It is possible to use FDRTCOPY to copy backups created by ABR's Application Backup (DUMP TYPE=APPL), but it is usually not useful. Application backup creates a Control File describing the backups, similar to the Archive Control File, but the usual process for application backup involves putting a backup of that Control File as the last file on the backup tape, in order to make that tape self-contained. FDRTCOPY can copy the backup tapes, and can update the Control File if it is still on disk, but it cannot update the backup of the Control File on the tape.

If you are using a variation of Application Backup which only creates a disk Control File, you may use FDRTCOPY on those backup; you will need to provide a ARCHIVE DD statement pointing to the control file, and you must use the ABRINDEX= operand to specify the high-level index used for those backups.

FDRTSEL contains features which can be used for copying Application Backups.

**SUPPORTED
DEVICES**

Since FDRTCOPY uses standard access methods, all tape devices and disk devices supported by your operating system are supported by FDRTCOPY. The input and output device types do not need to be the same, so FDRTCOPY can be used to copy backup data sets from disk to tape, or from one tape device type or density to another (such as 3400 tape reel to 3490E or Magstar (3590) tape cartridge). The output files may occupy more or fewer volumes than the original input.

**STORAGETEK
EXHPDM
SUPPORT**

FDRTCOPY supports input and output tapes which are managed by the ExHPDM software from StorageTek, as described in [Section 80.33](#). It is controlled by SUBSYS= operands on the TAPEIN, TAPEOUT, and TAPE2OUT DD statements.

60.11 FDRTCOPY JCL REQUIREMENTS

FDRTCOPY requires the following JCL to execute ([see Section 60.13](#) for examples):

EXEC STATEMENT	Specifies the program name (FDRTCOPY), and region requirement. A region of 4M is adequate for all functions.
STEPLIB DD STATEMENT	If required, specifies the load library in which FDRTCOPY resides. It must be an APF authorized library.
SYSPRINT DD STATEMENT	Specifies the output message data set. Normally a SYSOUT data set.
SYSUDUMP DD STATEMENT	Specifies the abend dump data set. Although not required, we strongly urge you to always include this DD statement, so that we can help you diagnose error conditions. Usually specifies a SYSOUT data set.
TAPEIN DD STATEMENT	<p>Specifies the first or only input tape or disk FDR backup data set. If the input is on tape and the ALLFILES or MAXFILES= are specified, the TAPEIN DD statement must specify all required information to OPEN the first file to be copied. At least DSN= and DISP= must be given; if the input data set is not cataloged then UNIT=, VOL=, and possibly LABEL= must also be given.</p> <p>If a multi-file, multi-volume ABR tape set is to be copied, normally only the first tape volume serial in the set must be given; FDRTCOPY will locate the additional volumes in the set in the ABR catalog. However, if some of the backups in the set are no longer cataloged, it may be necessary to specify all of the volume serials in the VOL= parameter.</p> <p>On tape, the copy may start at a file other than File 1 by specifying the file number in the LABEL=nn JCL parameter; however, the full data set name of the first file being copied must always be given.</p> <p>The ddname 'TAPEIN' may be overridden by the TAPEIN= parameter on the COPY statement.</p> <p>If an input tape is managed by the ExHPDM software product from StorageTek (see Section 80.33), you must indicate this by adding the SUBSYS= operand to the TAPEIN DD statement.</p>
TAPEOUT DD STATEMENT	<p>Specifies the primary output tape or disk to be created.</p> <p>The data set name specified will be overridden if COPYDSN=YES is specified on the COPY statement, or if the first input file is identified as an ABR-created backup. However, a proper data set name must be given on TAPEOUT or the system may treat this as a temporary data set.</p> <p>If multiple files are being copied, do not specify "DISP=(NEW,CATLG)"; instead use "DISP=(NEW,KEEP)" and use the CAT= parameter of the COPY statement to catalog output data sets.</p> <p>If more than 5 output tape volumes may be required, you must specify a volume count in the VOL= parameter, e.g., "VOL=(,,99)". Innovation recommends always specifying a volume count when the output is on tape.</p> <p>You may want to specify the RETPD= or EXPDT= JCL parameters to specify the expiration of the output file, but note that the EXP= operand in FDRTCOPY may override that expiration. If the output tapes already contain data, the copy may start beyond the existing files by specifying the starting file number in the LABEL=n JCL parameter.</p> <p>If the output is on disk, the DD statement can create the output file (with DISP=(NEW,KEEP) or (NEW,CATLG) and SPACE=) or can refer to an existing output file (DISP=OLD). However, remember that if COPYDSN=YES is specified (see the next section), the data set name copied from the input is the name that must exist on the output disk.</p>

60.11 CONTINUED

- TAPEOUT DD STATEMENT (continued)** TAPEOUT may be DUMMY; TAPEIN will still be read. This is useful to verify the readability of the input files, since FDRTCOPY will validate the contents and format of the backup file (Note: if the input file is compressed, specify COMPRESS=NONE for full validation).
- The DDNAME 'TAPEOUT' may be overridden by the TAPEOUT= parameter on the COPY statement.
- If an output tape will be managed by the ExHPDM software product from StorageTek ([see Section 80.33](#)), you must indicate this by adding the SUBSYS= operand to the TAPEOUT or TAPE2OUT DD statement.
- TAPE2OUT DD STATEMENT** (Optional) Specifies that a second output copy is to be created. All comments about 'TAPEOUT' above apply to 'TAPE2OUT'. The ddname 'TAPE2OUT' may be overridden by the TAPE2OUT= parameter on the COPY statement. If a TAPEOUT2 DD is present, it will be treated as an alias for TAPE2OUT. If the TAPE2OUT and TAPEOUT2 DDs are omitted, then the TAPEOUT file will be the only copy created.
- ARCHIVE DD STATEMENT** (Optional) Specifies the ABR Archive Control File. If ABR-created ARCHIVE backups are copied, this DD statement is required if the Control File is to be updated to point to the new copies unless the DYNARC parameter is specified on the COPY statement. COPYDSN=YES must be specified if the Archive Control File is to be updated for the first file copied. The ddname 'ARCHIVE' may be overridden by the ARCDD= parameter on the COPY statement.
- NOTE: If both the ARCHIVE DD statement and the DYNARC parameter are omitted when an ARCHIVE backup is copied, the Archive Control File will not be updated with the information about this copy. This backup can only be restored using the TAPEDD option of the RESTORE TYPE=ABR command.
- SYSIN DD STATEMENT** Specifies the control statement data set. Usually an input stream or DD * data set.

60.12 FDRTCOPY CONTROL STATEMENTS

COPY STATEMENT SYNTAX The COPY control statement must be present for FDRTCOPY to execute. However, all parameters are optional.

COPY	ABR= <u>YES</u> NO	,EXP= NO COPY DROP JCL
	,ABRCOPY= <u>SAME</u> FLIP n	,LASTAPE= dsname
	,ABRCOPY2= <u>SAME</u> FLIP n	MAXERR= nn
	,ABRINDEX= prefix	,NEWABRINDEX= prefix
	,ALLFILES	,PRINT= DSN
	,MAXFILES= nn	,PRINT= INFO
	,ARCBDSN= dsname	,SIMULATE
	,ARCDD= ddname	,SMSEXPIRE= YES
	,ARCEXPIRE= <u>UPDATE</u> KEEP NOT99000	,TAPEIN= ddname
	,CAT= <u>NO</u> YES RECAT	,TAPEOUT= ddname
	,CAT2= <u>NO</u> YES RECAT	,TAPE2OUT= ddname
	,COMPRESS= ALL COPY1 COPY2 NONE	,TAPERRCD= NO
	,COPYDSN= YES <u>NEW</u>	,TAPEXP= EXPDT RETPD
	,DROP= (EXPIRE,ARCBDSN,LASTAPE)	
	,DYNARC	
	,NODYNARC	

OPERANDS

ABR= **YES** – specifies that FDRTCOPY will look for and recognize ABR-created input files by the format of their data set names and invoke special processing for such files. [See Section 60.10](#) for details.
NO – specifies that FDRTCOPY will not recognize ABR formatted names on the input. NO should be used if you wish to copy ABR-created backups to non-ABR data set names.
 Default is YES.

ABRCOPY=
ABRCOPY2= If an input file is recognized as an ABR-created file these parameters control the copy number in the data set names of the equivalent output files.
 ABRCOPY= controls TAPEOUT, ABRCOPY2= is for TAPE2OUT (if present). The copy number is the digit in the second position in the last level of the data set name.
SAME – specifies that the copy number is not modified.
FLIP – indicates that it will be "flipped" to COPY 2 if COPY 1 is input, or to 1 if COPY 2 is input; this is especially useful for creating a second copy for offsite storage.
'n' is a single digit from 1 to 9 and sets the copy number. Incremental backups may have COPY 1 through 9 recorded in the ABR catalog, but archive backups may only have COPY 1 or COPY 2 recorded in the Archive Control File. These parameters may not be used with disk output.
 Default is SAME.

60.12 CONTINUED

- ABRINDEX=** Specifies the high level qualifier or prefix to be used in determining the fully-qualified data set name of ABR-created backup files. ABR files have a fixed naming convention (see [Section 52.05](#)) and FDRTCOPY supplies the hi-level index from the FDR/ABR Global Option Table.
But Application Backups (DUMP TYPE=APPL) may use a different prefix. To use FDRTCOPY to copy these application backups and recognize them as ABR files, you can specify ABRINDEX=userindex, where the 'userindex' is the first level qualifier of the application backup data set name.
- ALLFILES** – requests that all files on the input tape volumes be copied; the copy ends when the last file or last volume is detected (double tape mark or end of volume list).
MAXFILES=nn – requests the same processing except that no more than 'nn' files will be copied.
For files beyond the first, COPYDSN=YES is forced.
These parameters may not be used with disk input or output.
Default is to copy one file.
- ARCBDSN=** Specifies a data set name (up to 44 characters). The last 17 characters of the name will be compared to the dsname in tape labels on the input tapes to identify backup control files created by the ABR option ARCBACKUP=DSF (see [Section 51](#)). If the name ends in ".G0000V00", then the comparison will match on *any* GDG generation (any values in place of the zeros). The full dsname will be used to open the input file and the equivalent file on the output tape. If the DROP=ARCBDSN option is specified, these files will not be copied to the output tape. If omitted, no check for ARCBACKUP=DSF files will be done.
- ARCDD=** Specifies the DDNAME of the ABR Archive Control File. If ABR archive files are copied by FDRTCOPY, this DD statement must be present for the Archive Control File to be updated, unless DYNARC is specified.
Default is ARCHIVE.
- ARCEXPIRE=** When Archive or Application Backup files are being copied, this operand controls what expiration date will be recorded in the Control File for the output files created. See the "EXP=" operand for an explanation of how the expiration date of each output file is calculated.
UPDATE – specifies that the Control File will be updated with the new calculated expiration date.
KEEP – specifies that the expiration date currently recorded in the Control File will be left undisturbed. However, if the backup being recorded does not already exist (such as creating a new COPY 2 from a COPY 1), the calculated expiration will be recorded.
NOT99000 – if the expiration of the output file is 99000 (tape management catalog control), operate as if KEEP was specified, otherwise operate like UPDATE. This should be used if you keep your tapes under catalog control but record an actual expiration date in the Control File.
Default is UPDATE.

60.12 CONTINUED

CAT=
CAT2=

These two parameters control cataloging of data sets written by FDRTCOPY onto TAPEOUT and TAPE2OUT, respectively. Cataloging will be done only if the full data set name is known, i.e., if the first data set is being copied, if the original name was 16 characters or less, or if the data set was an ABR-created backup. When copying archive backups, the output file will be cataloged only if the equivalent input file was cataloged.

NO – inhibits cataloging.

YES – will cause the output data sets to be cataloged unless the name is already cataloged.

RECAT – catalogs even if it is already cataloged (updates the catalog).

NOTE: If both CAT= and CAT2= request cataloging, but the data set name on both output files is the same, only the data set on TAPEOUT will be cataloged.

Default is NO, except that if ABR=YES is in effect and the input file is identified as an ABR-created backup, the default is RECAT. If CAT= is specified but CAT2= is not, and there is a TAPE2OUT DD present, the default for CAT2= will be the value specified for CAT=. (see note at end of section)

COMPRESS=

ALL – specifies that all output backups will be in FDR compressed format.

COPY1 – causes only backups on TAPEOUT to be compressed.

COPY2 – causes only backups on TAPE2OUT to be compressed.

NONE – inhibits compression of any output.

Default is ALL if TAPEIN data is compressed, NONE if it is not.

NOTE: if the input file is compressed, and all output files are also compressed, then FDRTCOPY copies the data blocks without first decompressing and recompressing them, to reduce CPU time. However, in this case, some of the validation normally done by FDRTCOPY will be bypassed.

60.12 CONTINUED

COPYDSN= **YES** — specifies that for the first or only data set copied, the TAPEOUT/TAPE2OUT data set name will be the same as the input name (if an ABR-created file, the copy number may be modified if the ABRCOPY=/ABRCOPY2= parameters are present). The data set name is always copied for all files other than the first when using ALLFILES/MAXFILES=.

NEW — requests that FDRTCOPY assign a new name to the output data set. COPYDSN=NEW is only supported when FDRTCOPY is executed from FDRTEL, and the SMSEXPIRE=YES and ARCEDIT operands must also be specified in the FDRTEL input.

Without COPYDSN=NEW, FDRTEL will retain the original ABR name of the input backup data set. However, with ARCEDIT and SMSEXPIRE=YES you can selectively copy expired archived data sets from the input and leave other non-expired archived data sets in place in the input backup. Since both the input and output have the same data set names, only one of them can be cataloged. This causes an incompatibility with tape management "catalog control" and may cause one of the backups to expire prematurely if catalog control is used for these backups.

COPYDSN=NEW modifies the suffix of the output Archive backup data set according to the rules in [Section 51.01](#), searching for an appropriate suffix that is not currently cataloged. This allows you to keep the original input backup with its existing suffix, and create a new output backup with a new suffix for the data sets being moved.

WARNING: the Archive Control File (ACF) stores only one suffix to define the name of both copies of the backup of an archived data set; it simply varies the copy number in that suffix for COPY1 and COPY2. If you are using COPYDSN=NEW and the data sets being moved have 2 backup copies, you must provide both TAPEOUT and TAPE2OUT DD statements so that new backup data sets with the new suffix are created for both copies. Otherwise, the COPYx which was not re-created will no longer be accessible for these data sets.

Default is to use the data set name on the TAPEOUT/TAPE2OUT DD statement. However, if ABR=YES is in effect and the first input file is identified as an ABR-created backup, COPYDSN=YES is forced (see note at end of section).

DROP= Specifies conditions under which certain files on the input tapes will not be copied to the output.

EXPIRE — specifies input files which have reached their expiration dates (as recorded in the input tape label or disk DSCB) will not be copied (if the expiration on the input tape label is '99000', the data set will be considered expired if it is not cataloged, except for archive backup files which may not be cataloged).

ARCBDSN — if the ARCBDSN= operand described above is also specified, any ARCBACKUP=DSF files identified on the input tape will not be copied to the output tape.

LASTAPE — if the LASTAPE= operand described below is also specified, any LASTAPE files identified on the input tape will not be copied to the output tape.

You may specify any or all of the values of DROP=, placing them in parentheses if more than one is present, e.g., DROP=(EXPIRE, LASTAPE)

60.12 CONTINUED

DYNARC
NODYNARC

DYNARC – requests that the ABR Archive Control File named in the FDR/ABR Global Option Table be dynamically allocated for updating if Archive backups are copied.

NODYNARC – specifies that the Archive Control File will not be updated unless an ARCHIVE DD statement is present.

If NODYNARC is specified, no ARCHIVE DD is supplied, and Archive backups are being copied, a warning message will be issued to document that no Control File is being updated, which may make the backup copies unusable.

Default is NODYNARC. However, if ABR=YES is in effect, the input file is identified as an ABR-created Archive or Application Backup and no ARCHIVE DD statement is supplied, DYNARC is the default. (see note at end of section)

EXP=

Controls how FDRTCOPY will calculate the expiration date of output files. Also see the TAPEXP= operand.

NO – specifies that no expiration date processing will be done; the expiration date of output files will be obtained from the RETPD= or EXPDT= operands on the TAPEOUT/TAPE2OUT DD statements; if not specified there, an expiration date of 00.000 (already expired) will be assumed even though your operating system or tape management system may provide other defaults.

COPY – requests that the expiration date of each output file will be copied from the equivalent input file (from the tape label or disk DSCB).

DROP – is the same as COPY, except that input files whose expiration dates have been reached will not be copied. Although this option is still supported, we recommend using the DROP=EXPIRE operand in its place.

JCL – if a RETPD= or EXPDT= operand is specified on the TAPEOUT or TAPE2OUT DD statement, it will be honored. Otherwise, the expiration date will be copied from the equivalent input file (from the tape label or disk DSCB).

EXP= is supported for disk input. For disk output, the expiration will not be copied regardless of the EXP= operand.

NOTE: FDRTCOPY can copy only the expiration date recorded in the tape label or disk DSCB when the data set was created. Expirations that were changed in the Control File with the FDRARCH utility or changed in your tape management system are not known to FDRTCOPY.

Default is JCL.

LASTAPE=

Specifies a data set name (up to 44 characters). The last 17 characters of the name will be compared to the dsname in tape labels on the input tapes to identify dummy files created by the ABR option LASTAPE (see Sections 50 and 51). The full dsname will be used to open the input file and the equivalent file on the output tape. If the DROP=LASTAPE option is specified, these files will not be copied to the output tape. If omitted, no check for LASTAPE files will be done.

MAXERR=

Specifies the number of input file I/O errors and the number of block length errors that may occur before FDRTCOPY will abend the copy function (the count for these two conditions is maintained separately).

Default is 10 errors.

NEWABRINDEX= When ABR=YES is in effect, and the input file is identified as an ABR-created backup, this specifies that FDRTCOPY will replace the high-level index of the output data set name (the "abrindex") with the specified value (1 to 8 characters). This is intended for use with Application Backups (DUMP TYPE=APPL).

60.12 CONTINUED

PRINT=	DSN – specifies that FDRTCOPY is to list the names of the data sets and ICF VSAM components that are present in the backup file being copied. Default is not to list those names.
SIMULATE	Specified that FDRTCOPY is to only validate all operands on the COPY statement. No input or output files will be opened.
SMSEXPIRE=	Only honored when ABR=YES is in effect and the input file being copied is identified as an ABR-created Archive backup. This is primarily for use when FDRTCOPY is called by FDRTSEL. YES – causes the expiration date of the output file in the Archive Control File to be copied from the expiration of the existing COPY 2 of the same backup. This is designed to be used with Archive backups created with the ABR option SMSEXPIRE=YES, which sets the COPY 1 and COPY 2 expirations from SMS management class parameters. When the COPY 1 on disk (expiration set from LEVEL 1 DAYS NON USAGE) is moved to tape, it gets the same expiration as the COPY 2 (which was set to the final expiration of the archive). The actual expiration of the output backup file on tape should be set by EXPDT= or RETPD= to a value larger than any of the individual expirations (or perhaps EXPDT=99000 for catalog control).
TAPEIN=	Overrides the TAPEIN DDNAME.
TAPEOUT=	Overrides the TAPEOUT DDNAME.
TAPE2OUT=	Overrides the TAPE2OUT DDNAME.
TAPERRCD=	NO – Specifies that output tape I/O errors, if recovered by forcing a new output volume, will not cause any error indication at FDRTCOPY termination. Default is output tape I/O errors cause a U0888 abend at step termination, although the copy may complete successfully.

60.12 CONTINUED

TAPEXP=

When copying the expiration of an input tape to an output tape (see EXP=), FDRTCOPY may not be able to tell if the expiration date recorded in the tape label is a real expiration date, or was a keyword meaningful to a tape management system. This operand tells FDRTCOPY how to handle the copied expiration date.

EXPDT – specifies that FDRTCOPY will pass the input expiration date to the output tape as if EXPDT= was specified. For dates prior to 1997 or past 1999, these should be treated by your tape management system as a real date, but for dates in the range of 1998 to 1999 (and perhaps 1997 for CA-TLMS), they may be treated as keywords with special meanings.

RETPD – specifies that FDRTCOPY will pass all the input expiration dates to the output tape as if RETPD= was specified. This means that your tape management system will treat them all as real expiration dates (with the possible exception of xx000 expirations).

The default is RETPD unless the expiration date being passed is 99000 (tape management catalog control) when EXPDT is assumed. Since keyword dates other than 99000 are rarely used with ABR, this default is usually the correct choice. Override it only when you know that the default will not give correct results.

NOTE: The defaults for the CAT=, CAT2=, COPYDSN= and DYNARC operands have changed for FDRTCOPY V5.3 when ABR=YES is in effect and the input file is identified as an ABR-created backup. These defaults are different from V5.2 and earlier releases. However, Innovation feels that these changed defaults reflect the way that FDRTCOPY is used most of the time; i.e., that these operands are almost always specified when copying ABR backups, and that incorrect results can occur if they are omitted. These changes will prevent errors and reduce the required operands for many operations.

60.13 FDRTCOPY EXAMPLES

**COPY FDR OR
DSF BACKUP**

Copy a FDR or DSF backup data set. UNIT= and VOL=SER= are specified, but if the input backup was cataloged, they could be omitted; in that case the backup could be on tape or disk. If the backup occupied multiple tape volumes, VOL=SER=(vol1,vol2,,) would be specified. Since COPYDSN=YES is NOT specified, the data set name on the TAPEOUT DD statement will be used. TAPEOUT specifies a volume count in case more than 5 output tape volumes are required.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//TAPEIN      DD      DSN=FDR.BACKUP,DISP=OLD,
//              UNIT=TAPE,VOL=SER=B00023
//TAPEOUT     DD      DSN=FDR.COPY,UNIT=TAPE,
//              DISP=(NEW,KEEP),VOL=(, , 255)
//SYSIN       DD      *
              COPY
/*
```

**COPY FDR OR
DSF BACKUP
WITH EXHPDM**

This is the same as the previous example, except that the input and output backups are to be managed by the ExHPDM software from StorageTek, as indicated by the SUBSYS=SOV operands (SOV is the default subsystem name used by ExHPDM; your installation may have changed it). In this example both the input and output are ExHPDM, but you can copy from non-ExHPDM to ExHPDM, and vice versa.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//TAPEIN      DD      DSN=FDR.BACKUP,DISP=OLD,SUBSYS=SOV
//TAPEOUT     DD      DSN=FDR.COPY,UNIT=TAPE,DISP=(,CATLG),SUBSYS=SOV
//SYSIN       DD      *
              COPY
/*
```

**COPY ABR
BACKUP
CREATING
TWO COPIES**

Copy an ABR-created incremental backup tape, creating one copy which will replace the original, and a second copy which will be cataloged as a COPY2. ABRCOPY2=FLIP requests that the data sets on TAPE2OUT have their copy numbers "flipped", in this case from COPY1 to COPY2. CAT=RECAT and CAT2=RECAT is assumed so the COPY1 files will replace the original files in the catalog, and the COPY 2 files will be cataloged. EXP=JCL is assumed, but since there is no expiration or retention specified in the JCL, the original expiration dates of the input files are copied to the output tapes. The first backup to be copied is located through the ABR catalog, since UNIT= and VOL= are not specified. Since ALLFILES is specified, all of the backups on the tape set, starting with the specified backup, are copied. If the backups extended to multiple tape volumes, FDRTCOPY will locate the additional volumes through the ABR catalog and automatically call for them to be mounted.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//TAPEIN      DD      DSN=FDRABR.VSYSRES.C1000101,DISP=OLD
//TAPEOUT     DD      DSN=DUMM1,UNIT=3480,VOL=(, , 255),DISP=(NEW,KEEP)
//TAPE2OUT    DD      DSN=DUMM2,UNIT=3480,VOL=(, , 255),DISP=(NEW,KEEP)
//SYSIN       DD      *
              COPY    ALLFILES,ABRCOPY2=FLIP
/*
```

60.13 CONTINUED

**COPY AN ABR
ARCHIVE TAPE**

Copy a multi-volume ABR archive tape (which is probably also multi-file), creating a "COPY1" from a "COPY2" tape, and update the ABR Archive Control File (archived data sets cannot be easily restored from archive tapes unless the Control File is updated). ABRCOPY=FLIP causes the copy number in the input file to be changed from "2" to "1". LABEL=RETPD=90 causes the expiration date of the output file, in the tape labels, the Archive Control File, and the tape management system (if any) to be set to 90 days from today. The ARCHIVE DD statement is necessary for Archive Control File updating unless DYNARC is specified. Even though multiple volumes are to be copied, only the first volume serial needs to be specified on the TAPEIN DD statement; in most cases, FDRTCOPY can locate the other volumes required from the ABR catalog.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//ARCHIVE     DD      DSN=FDRABR.ARCHIVE,DISP=SHR
//TAPEIN      DD      DSN=FDRABR.VTS0001.B296120A,UNIT=CART,
//              VOL=SER=A00021,DISP=(OLD,KEEP)
//TAPEOUT     DD      DSN=DUMM1,UNIT=TAPE,LABEL=RETPD=90,
//              DISP=(NEW,KEEP),VOL=(, , 255)
//SYSIN       DD      *
              COPY    ALLFILES,ABRCOPY=FLIP
/*
```

**RECREATE A
DAMAGED
ABR TAPE**

An ABR tape (backup or archive) has been damaged or lost, so its duplicate copy is copied to recreate it. Even though only one volume of a multi-volume backup is damaged it is easiest to recreate the entire set of volumes. ABRCOPY=FLIP causes the copy number to "flip" from the good copy that of the copy being replaced. EXP=JCL is assumed but since no expiration is specified it copies the expiration date from the input files. This example shows reading a COPY2 archive to recreate a COPY1, but the input could be COPY2 or could be a backup. The DYNARC operand will cause the Archive Control File to be updated to point to the new copy, if the input is an archive tape.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//TAPEIN      DD      DSN=FDRABR.VTS0001.B290020A,UNIT=TAPE,
//              VOL=SER=A00079,DISP=(OLD,KEEP)
//TAPEOUT     DD      DSN=DUMM1,UNIT=TAPE,
//              DISP=(NEW,KEEP),VOL=(, , 99)
//SYSIN       DD      *
              COPY    ALLFILES,ABRCOPY=FLIP,DYNARC
/*
```

**COPY BACKUP
ON DISK TO
TAPE**

Copy a sequential disk FDR backup file to tape. The DISP= parameters shown will cause the disk file to be deleted and the tape copy to be cataloged if the copy is successful.

```
//FDRTCOPY    EXEC    PGM=FDRTCOPY,REGION=1M
//SYSPRINT    DD      SYSOUT=*
//SYSUDUMP    DD      SYSOUT=*
//TAPEIN      DD      DSN=FDR.BACKUP.VPROD03,
//              DISP=(OLD,DELETE,KEEP)
//TAPEOUT     DD      DSN=FDR.BACKUP.VPROD03,UNIT=TAPE,
//              DISP=(NEW,CATLG,KEEP)
//SYSIN       DD      *
              COPY
/*
```

60.13 CONTINUED

**AUTOMATED
COPY OF
ABR TAPES**

FDRTCOPY requires that the full data set name of the first data set to be copied be provided. The naming convention for ABR-created backup data sets causes the data set names to change with every ABR run, so a daily FDRTCOPY job would need to be manually updated every day. The following example is a technique to automate the copying of these tapes and avoid having to manually update JCL for each run.

FDRTCOPY can copy a file that was not created by FDR or ABR as long as it is empty, containing no records. These empty files will be copied but will not be cataloged, and will not generate an error. You can create an empty file of a known name as the first file on the ABR tape, using IEBGENER, and use that name in the FDRTCOPY JCL. In this example, the dummy file is a GDG, so the appropriate GDG index must be created beforehand; a non-GDG could also be used.

Note that if an ABR begins over on a new scratch tape (as after an abend of a backup task, or if the MAXFILES= value is reached) the files created after that point will NOT be included in the copy created by FDRTCOPY.

Changes to the Backup JOB.

```
//*          Create a Dummy data set as first file on tape
//DUMMYFIL   EXEC   PGM=IEBGENER
//SYSPRINT   DD     SYSOUT=*
//SYSIN      DD     DUMMY
//SYSUT1     DD     DUMMY,DCB=(RECFM=FB,BLKSIZE=80)
//SYSUT2     DD     DSN=FDRABR.DUMMY.BKUP(+1),
//            UNIT=TAPE,DISP=(,CATLG),VOL=(,RETAIN)
//*          Do an ABR Backup, starting at second file
//ABRSTEP     EXEC   PGM=FDRABR
//TAPE1      DD     DSN=FDRABR.DUMMY1,VOL=REF=*.DUMMYFIL.SYSUT2,
//            LABEL=(2,SL),DISP=(,KEEP)
//*          (remainder of ABR job step)
```

Changes to the FDRTCOPY JOB.

```
//*          Copy current backups relate to the dummy
//*          cataloged data set.
//TCOPY       EXEC   PGM=FDRTCOPY,REGION=1M
//SYSPRINT   DD     SYSOUT=*
//TAPEIN     DD     DISP=OLD,DSN=FDRABR.DUMMY.BKUP(0)
//TAPEOUT    DD     DSN=DUMM1,UNIT=TAPE,
//            DISP=(NEW,KEEP),VOL=(,,99)
//SYSIN      DD     *
//            COPY    ALLFILES,ABRCOPY=2
```


60.20 FDRTSEL INTRODUCTION

FUNCTION FDRTSEL automates the function of copying and moving ABR backup files.

For ARCHIVE backups:

- If you ARCHIVE with one copy to disk, FDRTSEL can move those disk archive backups to tape when they have reached (or are about to reach) their expiration dates, assigning the new tape copy a longer expiration. This allows you to create "level 1" archives on disk for quick recall, and later migrate them to "level 2" tape for longer retention. Criteria other than expiration date can be used when selecting backups for migration.
- FDRTSEL can be used to move archive backups from one media to another, such as disk to disk, tape to tape, or disk to tape.
- FDRTSEL can be used to recycle or consolidate archive tapes, by eliminating the archive backups of disk data sets that are no longer required (such as those that are expired, deleted, or recalled). It produces a set of output tapes which are smaller and use fewer volumes than the original, reducing the size of the archive tape library.
- FDRTSEL can create COPY 2 backups from COPY 1 if they were not created at ARCHIVE time, and can recreate one copy from the other if a copy is damaged or lost.

For Full-volume and incremental backups:

- Incremental and full-volume ABR backups can be copied. You can use this to create COPY 2 backups from COPY 1 if they were not created at backup time. Additionally, FDRTSEL can create extra off-site copies (3-9) and can recreate a copy from another if a copy is damaged or lost.
- FDRTSEL can be used to organize a generation's worth of backups for one or more disk volumes in order to minimize tape handling during disaster/recovery restores.

FUNCTIONAL DESCRIPTION

FDRTSEL internally invokes FDRTCOPY ([see Section 60.10](#)) to copy individual backups, but it automates selection and allocation of input files.

If copying or moving ARCHIVE files, FDRTSEL first invokes an ARCHIVE utility to scan the Archive Control File and select the archive backups that meet the user-specified selection criteria. It then allocates the first selected archive backup to the TAPEIN DDNAME and invokes FDRTCOPY to copy the archive backup to TAPEOUT (and TAPE2OUT), and to update the Archive Control File to reflect the new archive backup volumes. If the input backup is on disk, a successfully copied backup may optionally be scratched (and uncataloged if appropriate). The above process is repeated for each selected archive backup.

If copying ABR incremental or full volume backups, FDRTSEL first invokes an ABR catalog processor to select the backups that meet the user specified selection criteria. FDRTSEL then allocates the first selected ABR backup file to the TAPEIN DDNAME and invokes FDRTCOPY to copy the BACKUP file to TAPEOUT (and TAPE2OUT), and to catalog the created copies in the ABR catalog. The above process is repeated for each selected ABR incremental or full volume backup.

FEATURES

These features of FDRTSEL enhance its utility. Some of them were available in earlier releases of FDRTSEL, but may have been enhanced.

- ARCEDIT – when copying ARCHIVE backups, allows a partial copy of the Archive backup files by selecting or excluding certain of the original archived disk datasets (usually excluding expired datasets).
- LAST TAPE – allows you to add new files onto a tape created by a previous execution of FDRTSEL.
- DISK OUTPUT – allows FDRTSEL to copy backups from disk to disk. Can be used to move backups to new volumes, to convert the backups to a new device type (such as 3380 to 3390) or can be used in conjunction with ARCEDIT to reduce the size of the backups.
- CHECKPOINT/RESTART – allows long-running FDRTSEL jobs to be restarted after a failure or interruption.
- CONSOLE CONTROL – allows FDRTSEL to be interrupted gracefully (between backups) via a console command.

60.21 FDRTSEL JCL REQUIREMENTS

FDRTSEL requires the following JCL to execute:

EXEC STATEMENT Specifies the program name (FDRTSEL) and region requirement. The recommended region size is 4M or larger.

STEPLIB DD STATEMENT If required, specifies the load library in which FDRTSEL resides. It must be an APF authorized library.

SYSPRINT DD STATEMENT Specifies the output message data set. Normally a SYSOUT data set.

SYSUDUMP DD STATEMENT Specifies the abend dump data set. Although not required, we strongly urge you to always include this DD statement, so that we can help you diagnose error conditions. Usually specifies a SYSOUT data set.

TAPEIN DD STATEMENT Must **not** be present. It will be dynamically allocated by FDRTSEL.

TAPEOUT DD STATEMENT Specifies the destination for the primary or only output copy to be created by FDRTSEL. The ddname 'TAPEOUT' may be overridden by the TAPEOUT= parameter on the COPY statement.
SIMULATION: If CONTROL SIM is specified, you may omit TAPEOUT or specify it as:
//TAPEOUT DD DUMMY

TAPE OUTPUT: If outputting to tape or cartridge, specify:

UNIT= specify a generic (e.g., 3490) or esoteric (e.g., CART) name to allocate the type of tape drive desired. If you have sufficient tape drives available, specifying a unit count of 2 (e.g., UNIT=(3490,2)) may reduce elapsed time.

DISP=(NEW,KEEP) – do not specify CATLG since FDRTSEL handles cataloging of output files internally.

VOL= specify a volume count (e.g., VOL=(,,255)) to prevent FDRTSEL from abending if more than 5 tape volumes are required. If no volume serials are specified, FDRTSEL will call for scratch tapes; this is recommended; however, you may specify up to 255 tape volume serials.

LABEL= you may want to specify RETPD=nnn or EXPTD=yyddd to provide the expiration date of the backups. See the EXP= parameter in [Section 60.12](#) for details on handling of expirations.

If multiple backups are copied, FDRTSEL will create multiple files on the tape (or tape aggregate if more than one tape volume is used).

EXAMPLE: //TAPEOUT DD DSN=ABR1,UNIT=3490,DISP=(NEW,KEEP),
 // VOL=(, , 255), LABEL=EXPDT=99000

60.21 CONTINUED

**TAPEOUT DD
STATEMENT
(Continued)**

LAST TAPE OPTION: The LAST TAPE option of FDRTSEL allows you to add backup files to a tape created by a previous FDRTSEL step (even if that step is in another job and even if it was run on a previous day). This option is controlled entirely through JCL. To request LAST TAPE, the TAPEOUT DD is similar to that described above except that you specify:

DSN= any name that includes an index level of "LASTAPE". You may have multiple LASTAPE files for various purposes. This name will be cataloged to record the tape volume serial and file number where FDRTSEL is to start its output.

DISP=(MOD,KEEP) – this tells FDRTSEL to locate the LASTAPE file in the catalog, verify that the file exists on the output tape, and begin outputting to the tape at that point. If the name is not cataloged, FDRTSEL will call for a scratch tape and begin at file 1. Also, if you specify NEW instead of MOD, FDRTSEL will ignore the LASTAPE file and use scratch tapes (but it will still record the LASTAPE for future use).

VOL= volume serials should not be specified, but the volume count should be given.

```
EXAMPLE:      //TAPEOUT      DD      DSN=TECH.LASTAPE.COPY1,
              //              UNIT=3490,DISP=(MOD,KEEP),
              //              VOL=(, , 255),LABEL=EXPDT=99000
```

DISK OUTPUT: You may request that FDRTSEL copy the backups to disk. This is usually used with ARCHIVE backups, not full-volume or incremental backups. Disk output can only be requested if the input backup files are also on disk. In this case the TAPEOUT DD is used only to specify one or more disk volumes on which the backups will be placed; FDRTSEL will internally allocate the required backup files on those output disks.

To request disk output, specify the UNIT= device type, DISP=OLD, and VOL=SER= one or more disk volume serials, e.g.,

```
//TAPEOUT DD UNIT=3390,DISP=OLD,VOL=SER=(ARC001,ARC002,ARC003)
```

Alternately, you may catalog a dummy dataset to a set of output volumes and refer to that name in the JCL (remember that the name itself will not be used, only the volser list it points to), e.g.,

```
//TAPEOUT DD DISP=OLD,DSN=FDRABR.POOLDISK
```

NOTE: output disk volumes cannot be SMS-managed.

If an output tape will be managed by the ExHPDM software product from StorageTek ([see Section 80.33](#)), you must indicate this by adding the SUBSYS= operation to the TAPEOUT or TAPE2OUT DD statement. Archive backups cannot be copied to ExHPDM, but the volume backups and application backups can go to ExHPDM.

**TAPE2OUT DD
STATEMENT**

(Optional) Specifies the duplicate output tape or disk copy to be created. All comments about 'TAPEOUT' above apply to TAPE2OUT.

The ddname 'TAPE2OUT' may be overridden by the TAPE2OUT= parameter on the COPY statement. If a TAPEOUT2 DD is present, it will be treated as an alias for TAPE2OUT. If the TAPE2OUT and TAPEOUT2 DDs are omitted, then the TAPEOUT file will be the only copy created.

**ARCHIVE DD
STATEMENT**

Specifies the ABR Archive Control File to be used to identify archive backups to be copied when SELECT ARCHIVE is used. This ARCHIVE file will also be updated to reflect the results of the copy. If omitted, the DYNARC option will be assumed.

**TSELCKPT DD
STATEMENT**

(Optional) Specifies the Checkpoint recovery file. Its usage is described in [Section 60.27](#). If included, it should be allocated with JCL similar to:

```
//TSELCKPT DD DSN=PROD.TSEL.CHECKP,UNIT=DISK,SPACE=(CYL,(2,1)),
//          DISP=(MOD,CATLG)
```

This will create and catalog the file if it does not exist, and will use it if it does exist. If the file is new or empty, FDRTSEL will write checkpoint information to it. But if it is not empty at initialization, FDRTSEL will automatically perform a restart. If FDRTSEL completes successfully, it will rewrite the file as empty, ready for the next FDRTSEL execution.

60.21 CONTINUED

SYSIN DD STATEMENT Specifies the control statement data set. Usually an input stream or DD * data set.

RETURN CODES FDRTSEL will end with a return code as shown in this table:

0 (zero)	All Functions ended normally
4 or nn	All completed functions ended normally, but FDRTSEL was terminated prematurely because of the MAXFILES= operand (on the CONTROL statement) or because of a STOP console command. nn - Set by STOPCC=nn value on CONTROL stmt.
8	An error occurred (other than a FDRTCOPY error).
12	An FDRTCOPY error occurred

If no backups match the selection criteria, FDRTSEL will end with return code of 8, unless SELTERR=NO is specified. If SELTERR=NO, FDRTSEL will end with return code of 0 when there are no matching backups.

60.22 FDRTSEL CONTROL STATEMENTS

FDRTSEL accepts three primary control statements. These statements cannot be repeated. Each statement can only appear once within a given execution of FDRTSEL.

- **COPY or MOVE** – Specifies FDRTCOPY output processing options. MOVE also causes input disk files to be scratched.
- **SELECT** – Selects ABR files to be processed.
- **CONTROL** – Specifies additional FDRTSEL options. Can be used to indicate simulation mode. CONTROL is optional.

**SELECT
STATEMENT**

The SELECT statement identifies backup data sets to be copied or moved in ONLY ONE of three modes:

SELECT ARCHIVE is used to copy or move whole archive backup data sets. The SELECT ARCHIVE statement invokes FDRABRP to generate a list of backup data sets to be processed. All FDRABRP options available to PRINT ARCHIVE are honored by the SELECT ARCHIVE statement.

SELECT ARCDIT is identical to the SELECT ARCHIVE statement with one major difference. With SELECT ARCDIT, only the specifically requested individual data sets from the selected backup file get copied.

SELECT CATLG is used to copy ABR full-volume and incremental backup files. This statement creates a backup selection list by processing the ABR catalog.

60.23 FDRTSEL COPY/MOVE STATEMENT

**COPY/MOVE
STATEMENT**

The COPY or MOVE statement must be present. The MOVE statement can be used only with disk input; it causes the input backup file to be scratched (and uncataloged if appropriate) after it is successfully copied. COPY can be used with either disk or tape input, and it does not disturb the input backup; however, depending on your options, the output backup will may be cataloged in place of the input.

COPY	,ABRCOPY=<u>SAME</u> FLIP n	,MAXERR=nn
MOVE	,ABRCOPY2=<u>SAME</u> FLIP n	,NEWABRINDEX=prefix
	,ABRINDEX=prefix	,PRINT=DSN
	,ARCEXPURE=<u>UPDATE</u> KEEP NOT99000	,SIMULATE
	,CAT=<u>NO</u> YES RECAT	,SMSEXPIRE=YES
	,CAT2=<u>NO</u> YES RECAT	,TAPEOUT=ddname
	,COMPRESS=ALL COPY1 COPY2 NONE	,TAPE2OUT=ddname
	,COPYDSN=NEW	,TAPERRCD=NO
	,DYNARC	,TAPEXP=EXPDT RETPD
	,NODYNARC	
	,EXP=NO COPY DROP JCL	

All of the operands of COPY and MOVE shown above are also operands of the FDRTCOPY COPY statement; they will actually be passed to FDRTCOPY when it is invoked by FDRTSEL. Please see [Section 60.12](#) for the descriptions of these operands.

However, note these differences:

- Only those operands shown above are supported by FDRTSEL. Other FDRTCOPY operands documented in [Section 60.12](#) are not supported and should not be specified.
- If SELECT CATLG is used, the default for ABRCOPY= is set by the IFNOCOPY= or COPY= operands on that statement.
- If the SIM operand on the CONTROL statement is specified ([see Section 60.26](#)), then FDRTSEL will not even call FDRTCOPY. If that operand is not specified, but the SIMULATE operand is specified on the COPY/MOVE statement, FDRTSEL will allocate all of the required input files and invoke FDRTCOPY but FDRTCOPY will only validate its operands and exit.
- You should usually not specify the EXP= operand. For disk-to-disk MOVES, EXP=NO will result in incorrect operation (EXP=COPY is the default in that case). For all other operations, the default of EXP=JCL will usually result in the correct results.

60.24 FDRTSEL SELECT ARCHIVE/ARCEDIT STATEMENT

NOTE: Only one form of the SELECT statement can be used per execution of FDRTSEL. Only one SELECT ARCHIVE, SELECT ARCEDIT or SELECT CATLG statement can be present in a given job step.

**SELECT
ARCHIVE
STATEMENT**

The SELECT ARCHIVE statement is used to specify the criteria for selection of the archive backup data sets to be copied or moved. The most of the operands of the SELECT ARCHIVE statement are identical to the PRINT ARCHIVE statement of FDRABRP (See Section 53.05) with the following restrictions and additions:

- The IFONLYCOPY= operand limits the selection of the archive backups to those backups for which the other copy does not exist;
- The COPY= or IFONLYCOPY= operand must be specified to indicate which copy of the backup is to be processed.
- The BKDEVTYP= and BKVOL= operands may be used to select based on the backup type and/or volume serials.

FDRABRP will be internally invoked to process the Archive Control File and select archived data sets according to the operands you specify. If multiple selection criteria are specified, only those archived data sets which match all of them are selected. Note that if any one archived data set in a given archive backup is selected, the entire backup will be processed.

**SELECT
ARCEDIT
STATEMENT**

The SELECT ARCEDIT statement can be used instead of the SELECT ARCHIVE statement. The ARCEDIT function operates exactly the same as the SELECT ARCHIVE, with the only exception being that it "edits" the archive backup file written to TAPEOUT to include only the archived data sets that were selected from the Archive Control File. ALL other keywords and functions are identical to the SELECT ARCHIVE statement.

SELECT	ARCHIVE ARCEDIT	,PDATE=yyddd yyyyddd
S	,ADATE=yyddd yyyyddd	,PDAYS=nnnn
	,BKDEVTYP=<u>DISK</u> TAPE ANY	,RESTORED=YES NO
	,BKVOL=(vvvvvv,...,vvvvvv)	,SDATE=yyddd yyyyddd
	,COPY=1 2	,SDAYS=nnnn
	,DELETE=YES NO	,SELTERR=<u>YES</u> NO
	,EXPDATE=yyddd yyyyddd	,VOL=(vvvvvv,...,vvvvvv)
	,EXPDAYS=nnnn	,XDAYS=nnn
	,EXPIRE=YES NO	,XVOL=(vvvvvv,...,vvvvvv)
	,IFONLYCOPY=1 2	

60.24 CONTINUED

**SELECT
ARCHIVE
OPERANDS****ADATE=**

Specifies that the data set was archived. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123=2003.123). FDRTSEL will only SELECT the data sets which match this date.

The default is that the date is not checked.

BKDEVTYPE=

DISK – Specifies the selection is limited to only those archive backups that reside on disk.

TAPE – Specifies the selection is limited to only those archive backups that reside on a tape media. This is determined by the presence of a file number for the backup COPY1 or COPY2 information.

ANY – Either disk or tape backups may be selected.

The default is DISK.

BKVOL=

Specifies one or more volume serials or volser prefixes of archive backup volumes; only those archive entries indicated as being archived to one of those volumes will be selected. These may be either tape or disk volume serials. To specify a volser prefix, follow the prefix with an asterisk (*). For example, BKVOL=ARCH* selects all volumes whose serials begin with ARCH. Multiple volume serial numbers may be specified if entered in parentheses, e.g., BKVOL=(ARCH*,ARC001). Note that the backup volume is only compared for the specific COPY= being selected.

The default is that the backup volume serial will not participate in data selection.

Note: The BKVOLG= operand, used in previous releases to select a backup volume prefix, is no longer accepted. Use BKVOL=prfx* instead.

COPY=

Specifies which backup copy (1 or 2) is to be copied or moved. If COPY=2 is specified backups which do not have a COPY 2 will be bypassed.

FDRTSEL requires the COPY= or IFONLYCOPY= keyword to be specified. Use COPY=n to unconditionally create a copy of this FDR file.

DELETE=

YES selects only the archived data sets that have been flagged for deletion and are subject to removal from the Archive Control File.

NO selects only the archived data sets that have not been set for deletion.

The default is that delete status will not participate in data selection.

60.24 CONTINUED

EXPDATE= EXPDATE= specifies an expiration date. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123 = 2003.123).

EXPDAYS=

EXPDAYS= specifies a value in days used to calculate a future expiration date (today plus "nnnn").

In addition to the equal sign (select data sets with a single expiration), you can use a variety of operators to select a range of expiration dates. Supported operands are:

=	or	.EQ.	equal
¬ or ¬=	or	.NE.	not equal
<	or	.LT.	less than
>	or	.GT.	greater than
<=	or	.LE.	less than or equal to
>=	or	.GE.	greater than or equal to

If you use an operand other than equal, you can specify either operand **twice** on a control statement, to define a range of expiration dates.

For example:

EXPDAYS>30,EXPDAYS.LT.90

will select data sets with an expiration more than 30 and less than 90 days from today.

You can select multiple expiration dates or ranges by specifying each on a separate FDRTSEL control statement. If the selection criteria match on ANY of the expire ranges then the file is selected for copying

This command operates similarly to the XDAYS operand, however EXPDAYS and EXPDATE allow you to create an explicit expiration date range, or sets of ranges. The EXPDAYS and EXPDATE operands are NOT dependent on the EXPIRE=YES operand.(XDAYS= is).

If XDAYS is also specified, this criteria check is performed after the XDAYS processing. EXPDAYS and EXPDATE can both be specified on the same control card for a maximum of 2 values for each keyword.

The default is no expiration date checking will be done.

IFONLYCOPY= The IFONLYCOPY= operand limits the selection to those backups for which the other copy does not exist. To create a COPY 2 of any selected backup that does not already have a COPY 2, specify IFONLYCOPY=1.

FDRTSEL requires that either the COPY= or the IFONLYCOPY= keyword must be specified. Use IFONLYCOPY=n to create copies for which the other copy does not exist.

PDATE= Specifies that ONLY data sets archived on or BEFORE the date specified will be selected. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123=2003.123). This option can reduce execution time if many data sets have been archived.

60.24 CONTINUED

PDAYS=	<p>Specifies a value in days which is subtracted from today's date to calculate a prior Julian date. This date is used as described above by PDATE operand. This option can reduce execution time if many data sets have been archived.</p> <p>If neither PDATE= nor PDAYS= is specified, the default is that the entire Archive Control File will be searched.</p>
RECALL=	<p>YES selects ONLY the data sets that were archived by FDRABR with the RECALL option.</p> <p>NO selects only the data sets that were not been archived with RECALL.</p> <p>The default is that recall status will not participate in data selection.</p>
SDATE=	<p>Specifies that ONLY data sets archived on or after the date specified will be selected. The date is specified in Julian format with a 2-digit year (yyddd) or a 4-digit year (yyyyddd). If the 2-digit year is used, year numbers less than 70 will be assumed to be in the 21st Century (e.g., 03123=2003.123). This option can reduce execution time if many data sets have been archived.</p> <p>The default is that the entire Archive Control File will be searched.</p>
SDAYS=	<p>Specifies a value in days which is subtracted from today's date to calculate a prior Julian date. This date is used as described above by SDATE operand. This options can reduce execution time if many data sets have been archived.</p> <p>If neither SDATE= nor SDAYS= is specified, the default is that the entire Archive Control File will be searched.</p>
SELTERR=	<p>YES — Specifies that a condition code of 12 will be set if there are no archived data sets that match the selection criteria on the SELECT.</p> <p>NO — Specifies that a condition code of 0 will be set if there are no archived data sets matching the selection criteria on the SELECT and there are no other errors.</p> <p>The default is YES.</p>
VOL=	<p>Specifies one or more DASD volume serials or volser prefixes; only those archive entries indicated as being archived from one of those volumes will be selected. To specify a volser prefix, follow the prefix with an asterisk (*). For example, VOL=TSO* selects all volumes whose serials begin with TSO. Multiple volume serial numbers may be specified if entered in parentheses, e.g., VOL=(IPL001,MVS230,WORK03)</p> <p>The default is that the DASD volume serial will not participate in data selection.</p> <p>The VOLG= operand, used in previous releases to select a DASD volume prefix, is still accepted.</p>
XDAYS=	<p>Specifies a value in days which is added to today's date to calculate a future expiration date.</p> <p>XDAYS= is used with EXPIRE=YES to select data sets that will expire in the next nnn days.</p> <p>The default is 10 days, but is ignored unless EXPIRE=YES is coded.</p>

60.24 CONTINUED**XVOL=**

Specifies the DASD volume serial number from which the data set(s) were archived to EXCLUDE from the selection criteria, if other selection criteria does select this volume.

XVOL= can specify prefix values by using an asterisk as the last character. Multiple volume serials can be specified if entered in parenthesis, e.g.,

XVOL=(IPL001,MVS230,WORK*)

The default is that the volume serial will not participate in the selection criteria.

60.25 FDRTSEL SELECT CATLG STATEMENT

**SELECT CATLG
STATEMENT**

The SELECT CATLG statement is used to select ABR full-volume and incremental backups to be copied. The syntax of the SELECT CATLG statement is similar to the PRINT CATLG statement of FDRABRP (see [Section 53.06](#)) with the following restrictions and additions:

- The IFNOCOPY=n keyword limits the copies to be made to ABR backups which do not already have a copy "n".
- You can specify which copy to use as input by SOURCE=n.
- MAXGEN defaults to 1, selecting only the most recent generation, all cycles.

SELECT	CATLG	,MAXCYC=nn
S	,BKDEVTYP=DISK TAPE <u>ANY</u>	,MAXGEN=nnnn
	,COPY=n	,SELTERR=<u>YES</u> NO
	,CYCLES=nn	,SOURCE=n
	,GEN=nnnn	,TYPE=ABR FDR
	,IFNOCOPY=n	,VOL=(vvvvvv,...,vvvvvv)

**SELECT
CATLG
OPERANDS**

BKDEVTYP=	<p>DISK – Specifies the selection is limited to only those backups that reside on disk.</p> <p>TAPE – Specifies the selection is limited to only those backups that reside on a tape media. This is determined by the presence of a file number in the ABR catalog entry.</p> <p>ANY – Either disk or tape backups may be selected.</p> <p>The default is ANY.</p>
COPY=	<p>Specifies the copy number (1 to 9) you want to create even if that copy already exists. This value is passed to FDRTCOPY as ABRCOPY=n. The default (if the IFNOCOPY= keyword is not specified) is COPY=2.</p> <p>Note that if your JCL includes a TAPE2OUT DD statement (to create a second output copy), you must give the copy number of that copy by including the ABRCOPY2=n operand on the COPY or MOVE statement.</p>
CYCLES=	<p>Specifies the number of cycles (across generations) from and including the most current backup to be copied. This keyword can be used to select the nn most recent backups, independent of their generation.</p> <p>This keyword must not be specified with MAXGEN= or MAXCYC=.</p>
GEN=	<p>Specifies a specific generation number to be used for selection. If backups for this specific generation exist for the volume, they will be selected.</p> <p>This keyword can not be used with MAXGEN= or CYCLES=.</p>
IFNOCOPY=	<p>The IFNOCOPY= operand limits the copying selection to only backups for which you don't already have a copy "n" (1 to 9). For example, to create a COPY 2 of any backup for which the COPY 2 does not exist, specify, IFNOCOPY=2. This value is passed to FDRTCOPY as ABRCOPY=n.</p>

60.25 CONTINUED

- MAXCYC=** Specifies the number of cycles to be selected from each generation of a given disk volume. The order of selection will be most recent to least recent.
- The default is that ALL cycles within a generation will participate in the selection.
- MAXGEN=** Specifies the number of generations to be copied for each disk volume. Starting from the most current to the least current (highest number to lowest).
- The default is only the current generation will be used for selection of backups.
- SELTERR=** **YES** – Specifies that a condition code of 12 will be set if there are no cataloged backups that match the selection criteria on the SELECT.
NO – Specifies that a condition code of 0 will be set if there are no cataloged backups matching the selection criteria on the SELECT and there are no other errors.
- The default is YES.
- SOURCE=** Specifies the copy number (1 to 9) of the ABR backup to be used as input. All output copies will be generated from the SOURCE= copy. This copy must be cataloged in the ABR catalog.
- The default is to use the Copy1 (C1) file.
- TYPE=** **FDR** – limits the selection to only FDR full volume backups. (cycle = 00)
ABR – limits the selection to only ABR incremental backups. (cycle > 00)
- The default is to use both full volume and all incremental backups in the selection process.
- VOL=** Specifies one or more DASD volume serials or volser prefixes; selection will be limited to backups of those volumes. To specify a volser prefix, follow the prefix with an asterisk (*). For example, VOL=(MVS*,TSO*,PROD01) selects all volumes whose serials begin with MVS and TSO plus the volume PROD01.
- The default is that all cataloged backups are eligible for selection.

60.26 FDRTSEL CONTROL STATEMENT

CONTROL STATEMENT The CONTROL statement is optional. It specifies FDRTSEL processing options.

CONTROL	ABRINDEX= prefix ,ARCBACKUP= DSF ,ARCB1DSN= dsname ,ARCB2DSN= dsname ,MAXFILES= nnnn ,NOPOOL ,REBUILD ,SIM	,STACK= YES NO VOLUME nnnn ,STACK2= YES NO VOLUME nnnn ,STOPCC= nnn ,TMS NOTMS ,UNIT= 1 2 ,UNITNAME= (generic,newunit) ,UPPER
----------------	---	--

CONTROL OPERANDS	ABRINDEX= ARCBACKUP= DSF ARCB1DSN= ARCB2DSN= MAXFILES= NOPOOL	<p>Specifies the ABR prefix (first index level) of the ABR ARCHIVE and BACKUP files to be copied. This may be used with application backups (DUMP TYPE=APPL, see Section 51).</p> <p>The default is to use the ABRINDEX in the ABR global options table (usually "FDRABR").</p> <p>When copying Archive or Application Backup files, directs FDRTSEL to invoke FDRDSF to DUMP the Archive Control File used in this step as the last file on the tape once all files have been copied. This operates the same as the ARCBACKUP=DSF option of ABR.</p> <p>If the ARCB1DSN= and/or ARCB2DSN= operands are not specified, FDRTSEL will name these files on the tape by changing the index level "ARCHIVE" in the Archive Control File data set name to "ARCBKUP" for TAPEOUT and "ARCBKU2" for TAPE2OUT.</p> <p>Specifies the data set names to be used on TAPEOUT (ARCB1DSN=) and TAPE2OUT (ARCB2DSN=) for the backup of the Archive Control File if the ARCBACKUP=DSF operand is also specified. The default names are described under ARCBACKUP= above.</p> <p>Specifies the maximum number (1 to 9999) of ARCHIVE or BACKUP files to be copied in this FDRTSEL step. Once this limit is reached FDRTSEL terminates. Files not processed may be selected in a subsequent run.</p> <p>The default is that all selected files will be copied.</p> <p>For use when TAPEOUT is on disk. This option directs FDRTSEL to allocate the new copy of the backup file on the disk volumes in the order as specified in the TAPEOUT DD statement. By default, it will reorder the volumes in the list according to the amount of free space on each volume.</p>
-------------------------	--	--

60.26 CONTINUED

REBUILD	<p>Applies only to SELECT ARCHIVE/ARCEDIT functions.</p> <p>Directs FDRTSEL to select the backups for copying based on the COPY=n or IFNOCOPY=n operand, however it uses the other copy to copy from. This allows you to "REBUILD" a version of a set of backups from the other copy.</p> <p>NOTE: REBUILD also changes the default passed to FDRTCOPY to ABRCOPY=FLIP (i.e. to create the opposite copy of the version being read).</p>
SIM	<p>Invokes a simulation mode, and generates a report detailing the archive or backup files which would be copied. This is recommended for testing FDRTSEL options. Note that FDRTCOPY will not be invoked, and many options on the COPY/MOVE statement will not be validated (use the SIMULATE option on COPY/MOVE without the SIM option on CONTROL to test COPY/MOVE options).</p>
STACK= STACK2=	<p>Controls how output files are stacked on TAPEOUT (STACK=) and TAPE2OUT (STACK2=), i.e., whether multi-file output tapes will be created.</p> <p>YES – Stack up to 255 files.</p> <p>NO – Create only one file per output tape.</p> <p>VOLUME – Stack all the selected backups related to the same disk volume on one set of output tapes.</p> <p>nnnn – Stack up to nnnn files (1 to 4095).</p> <p>If the stack limit is reached, FDRTSEL will call for a fresh scratch tape (or the next tape in the TAPEOUT volume list if supplied) and start with file sequence 1.</p> <p>The default is YES.</p>
STOPCC=	<p>When premature termination of FDRTSEL is requested via the operator command "P jobname" or "F jobname,STOP" (see Section 60.27), FDRTSEL shuts down processing files at the completion of the current file. STOPCC= specifies the step return code to be set to indicate that this has occurred. The default is 4.</p>
TMS NOTMS	<p>Specifies the Tape Management System (TMS) option for use with the LAST TAPE option (see Section 60.27).</p> <p>NOTMS causes FDRTSEL to overwrite the LASTAPE file on the previous tape when adding new files to the tape.</p> <p>TMS causes it to add new files after the previous LASTAPE file for compatibility with tape management.</p> <p>The default is taken from the TMS option in the FDR/ABR Global Options table (ISPF panel A.I.4.4).</p>
UNIT=	<p>Directs FDRTSEL to request 1 or 2 tape devices when allocating input backups to be copied; 2 units may improve performance by minimizing waits for tape mounts and rewinds.</p>
UNITNAME-	<p>The first name is the genericc unitname that FDRTSEL will normally use to allocate the TAPEIN DD statement for a given backup on a given device, such as 3480, 3480X, 3490, or 3590-1. The second name is the user-supplied replacement esoteric or generic device name to use instead of the first name.</p> <p>Example: UNITNAME=(3490 , NONS I L O)</p> <p>This keyword can be repeated up to 15 times for a total of 16 UNITNAME device pairs.</p>
UPPER	<p>By default, FDRTSEL issues messages in upper and lower case print characters. If you require that all messages be printed in UPPER case only, specify this option.</p>

60.27 FDRTSEL FEATURES AND CONSIDERATIONS

This section describes some special features of FDRTSEL, and considerations for using the features of FDRTSEL.

**LAST TAPE
SUPPORT**

FDRTSEL has LAST TAPE support, similar to the LAST TAPE option in FDRABR (described in [Section 51](#)). LAST TAPE support allows you to add new files onto a tape created by a previous execution of FDRTSEL in a previous step or job, even if that job was run on an earlier day. LAST TAPE works only with tape output.

To use LAST TAPE, change the data set name in the TAPEOUT and/or TAPE2OUT DD statements to include a qualifier of "LASTAPE" anywhere in the name, e.g., "DSN=PROD.ARCHIVE.LASTAPE" or "DSN=TECH.LASTAPE.DAILY". You must also change the disposition to DISP=(MOD,KEEP).

If FDRTSEL finds a LASTAPE data set name, it invokes special processing to keep track of the last tape used for output so that it can add files to that tape:

- at the *end* of processing, after all files have been written to the output tape, FDRTSEL will write a dummy (empty) file with the LASTAPE dsname as the last file on that output tape, and catalogs it for reference.
- at the *start* of processing, it does a LOCATE on the name to get the tape volume and file number from the catalog. It then uncatalogs the backup (in case of abend) and opens the file to verify that it is still on the tape. If the LOCATE or the OPEN fails, it assumes that the last tape is not usable and starts outputting to a fresh scratch tape.
- if the LASTAPE file is successfully opened, FDRTSEL will add new files to the tape. It will overwrite the LASTAPE file unless the TMS (tape management) option is specified on the CONTROL statement or is set in the FDR/ABR Global Options Table (ISPF panel A.I.4.4), in which case it starts just beyond the LASTAPE file.

If you change the DISP to (NEW,KEEP), it will uncatalog the LASTAPE file (if it exists) but will start outputting to a fresh scratch tape without trying to open the file, allowing you to specify when to add files to the LASTAPE and when to start fresh by simply changing the DISP. Alternately, you can uncatalog the LASTAPE file and the next FDRTSEL run using that name will use a scratch output.

**DISK TO DISK
OPERATIONS**

When the TAPEOUT or TAPE2OUT DD statement points to a DISK volume, FDRTSEL will use this DD statement only to identify what volumes to use to receive the output files; data set name and other parameters on the DD are ignored. Internally, it will determine the available free space on all disk volumes identified by the TAPEOUT DD statement and will dynamically allocate a work DD statement with the disk volumes in the order of largest available free space. FDRTSEL recomputes the available free space and resets the order of the output volumes after each archive backup file is successfully copied (the NOPOOL operand on the CONTROL statement will override this sorting and use the volume in the order specified). [Section 60.21](#) has more information on the DD statement used with disk output.

If the output copy number (specified by ABRCOPY= or ABRCOPY2=) is different from input backup, FDRTSEL will simply COPY the requested backup files to the new output volumes. But if the output copy number is the same as the input, then FDRTSEL will copy the selected files to the new volsers with a copy number of "0". If the copy is successful, FDRTSEL will SCRATCH the input file from the disk volume, and then RENAME the "0" copy to the original copy number. The Archive Control File is updated appropriately.

Disk to Disk functions default to EXP=COPY to preserve the expiration date on disk. The input files selected by FDRTSEL for disk to disk must be cataloged and **you must not specify CAT=NO**. If CAT=NO is specified, this can cause the source TAPEIN file to be scratched and the file it was copied to will not be renamed properly.

Just like FDRABR pooldisk functions, FDRTSEL allows you to create a pool of volumes to be used in selection for the output of the disk to disk copy. You can define disk pools for the TAPEOUT dd statement by creating catalog entries pointing to the volsers or just allocated them via the JCL.

60.27 CONTINUED**CHECKPOINT/
RESTART**

FDRTSEL includes an option to checkpoint all of its processing, and to restart if it is interrupted. This is useful for long FDRTSEL jobs to avoid repeating operations already completed.

If the optional TSELCKPT DD statement (described in the next section) is present, FDRTSEL will record on this data set:

- all control card information (all options and selection criteria)
- a list of all backups selected (to be processed)
- if ARCEDIT is used, a list of all data sets selected from within each backup
- a completion record for each backup successfully processed.

During initialization, if the TSELCKPT data set is empty or new, FDRTSEL assumes normal operation and initializes the data set with the information listed above. If FDRTSEL completes normally, it rewrites the checkpoint file as empty, ready for reuse by the next FDRTSEL execution.

But If FDRTSEL is stopped or shutdown before completion due toabend, MAXFILES, operator cancellation or the console commands described below, the checkpoint file contains all of the information to restart FDRTSEL after the last successfully copied file. When the FDRTSEL job is resubmitted, and the TSELCKPT DD points to the existing checkpoint file, FDRTSEL will detect that the file is NOT empty and will begin recovery processing automatically. It loads the selection criteria from the checkpoint file, ignoring the criteria in the jobstream, and processes the entire selection list again. Before copying a backup, FDRTSEL checks to see if there is a completion record in the checkpoint for that backup. If so, it displays the original completion date and time and bypasses the backup. Any backup not recorded as completed will be processed.

**OPERATOR
COMMANDS**

FDRTSEL has an operator communications function which allows you to STOP a currently running FDRTSEL job and display the STATUS of an FDRTSEL JOB.

You can direct FDRTSEL to terminate after completing processing on the current input file by issuing either the console MODIFY (F) or STOP (P) command, substituting the job name of the FDRTSEL job:

F jobname,STOP

P jobname

To request that FDRTSEL display on the console number of files selected for copying and the current file being copied, issue the MODIFY (F) command:

F jobname,STATUS

60.27 CONTINUED

**CONSOLI-
DATING
ARCHIVE
TAPES**

An important function of FDRTSEL is the maintenance of your library of archived data on tape. Since archived data sets are usually kept for a long time, a year or more, the number of tape volumes devoted to archived data may become large. However, depending on how you create and managed your archived data, much of the data on those tapes may become obsolete, resulting in waste. This can occur if:

- you have data sets with varying expiration dates on the same tapes. This is especially likely if you use the SMSEXPIRE=YES function ([see Section 51](#)).
- you discard archived data sets which are no longer cataloged (such as GDGs).
- you don't keep the archived copies of data set after they have been recalled.

In these cases, you will want to periodically copy all or part of your archive library in order to discard the tape data that is no longer needed and consolidate the remainder onto a smaller set of tape volumes.

When you run FDRTSEL with SELECT ARCHIVE, it will select only those archive backup files which still have entries in the Archive Control File, so other backups on the input tapes will not be copied and the output tapes will contain less data.

SELECT ARCEDIT works much the same way, except that while copying the remaining backup files, the DASD data sets in those backups will be "edited" so that only those data sets which had entries selected from the Archive Control File will be copied; data sets which have been purged from the control file (or which were not selected) will be discarded, further reducing the size of those backup files which do remain.

The most straightforward way of consolidating archive tapes is to maintain your Archive Control File by periodically running the FDRARCH utility ([see Section 51.50](#)) with appropriate parameters to delete the entries for obsolete data sets, such as expired backups, restored data sets, uncataloged data sets, or whatever other options your installation's policy requires. Then run FDRTSEL to copy only the archive data for the data sets that remain in the control file. See the "Tape Consolidation" example in [Section 60.28](#).

NOTE: FDRTSEL has no interface to tape management systems. Once FDRTSEL has created a replacement set of consolidated Archive backup tapes, it is your responsibility to expire the original tapes to free them up.

60.28 FDRTSEL ARCHIVE EXAMPLES

**EXPIRED
DISK-TO-TAPE**

Select the COPY 1 ARCHIVE backups on disk (COPY=1 BKDEVTYP=DISK) that have expired or will expire within 3 days (EXPIRE=YES XDAYS=3) and move them to tape. Copy the selected ARCHIVE backups creating 2 tape copies (COPY 1 and COPY 2) which will be retained for 2 years, then scratch the archive backup from disk. The CONTROL statement indicates that a maximum of 10 files will be created on COPY 1 (TAPEOUT) before requesting a new tape, but up to 255 files will be stacked on COPY 2 (TAPE2OUT) and a maximum of 50 archive backups will be copied in this run. The CONTROL statement can be omitted to remove these restrictions.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR , DSN=FDRA BR . ARCHIVE
//TAPEOUT       DD     DISP=( , KEEP) , DSN=DUMMY1 , UNIT=( TAPE , 2) ,
//              VOL=( , , 255) , LABEL=RETPD=730
//TAPE2OUT      DD     DISP=( , KEEP) , DSN=DUMMY2 , UNIT=( TAPE , 2) ,
//              VOL=( , , 255) , LABEL=RETPD=730
//SYSIN         DD     *
MOVE           ABRCOPY2=2
SELECT         ARCHIVE , COPY=1 , BKDEVTYP=DISK , EXPIRE=YES , XDAYS=3
CONTROL        STACK=10 , STACK2=YES , MAXFILES=50
/*
```

**CREATE TAPE
COPY OF DISK
ARCHIVES**

Select the archive backups on disk (BKDEVTYP=DISK is the default) for which a second copy does not exist (IFONLYCOPY=1). Create a COPY 2 on tape with 1-year retention. This allows you to do archiving to disk only for quicker execution, and create the second copy on tape at a later time.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR , DSN=FDRA BR . ARCHIVE
//TAPEOUT       DD     DISP=( , KEEP) , DSN=DUMMY1 , UNIT=3490 ,
//              VOL=( , , 255) , LABEL=RETPD=365
//SYSIN         DD     *
COPY           ABRCOPY=2
SELECT         ARCHIVE , IFONLYCOPY=1
/*
```

**UNRESTORED
DISK-TO-TAPE**

Select the COPY 1 ARCHIVE backups on disk (BKDEVTYP=DISK is the default) created more than 60 days ago (PDAYS=60) and create a COPY 2 on tape, retaining only the archived DASD data sets which have not been recalled (ARCEDIT and RESTORED=NO). Archived data sets which have been removed from the Archive Control File by the FDRARCH utility will also be omitted from the copy.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4096K
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR , DSN=FDRA BR . ARCHIVE
//TAPEOUT       DD     DISP=( , KEEP) , DSN=DUMMY1 , UNIT=3480 ,
//              LABEL=RETPD=365
//SYSIN         DD     *
COPY           ABRCOPY=2
SELECT         ARCEDIT , COPY=1 , RESTORED=NO , PDAYS=60
/*
```

60.28 CONTINUED

**TAPE
CONSOLI-
DATION**

Select the COPY 1 ARCHIVE backups on tape for a full year (1995 in this example) and create new COPY 1 tapes, copying only the archived DASD data sets which are still recorded in the Archive Control File. Any tape files which no longer have any recorded data sets in the Control File will be completely dropped (not copied), while the use of ARCEDIT will insure that the copies of the remaining files will contain only the recorded data sets. This will result in a much smaller set of tapes. The original tapes will have to be manually expired in your tape management system. The use of the TSELCKPT data set allows the FDRTSEL job to be resubmitted and restarted if it fails or must be prematurely terminated (the second step deletes the checkpoint dataset only if the FDRTSEL step completes normally).

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TSELCKPT      DD     DSN=TECH.FDRTSEL.CKPT, UNIT=DISK, DISP=(MOD,CATLG),
//              SPACE=(CYL,(2,1))
//TAPEOUT       DD     DISP=(,KEEP), DSN=DUMMY1, UNIT=3490,
//              LABEL=RETPD=365
//SYSIN         DD     *
                COPY   DYNARC
                SELECT  ARCEDIT, COPY=1, BKDEVTYP=TAPE,
                        SDATE=95000, PDATE=96000
/*
//DELCKPT      EXEC   PGM=IEFBR14, COND=(0,NE,FDRTSEL)
//TSELCKPT      DD     DSN=TECH.FDRTSEL.CKPT, DISP=(OLD,DELETE)
```

**REPORT ON
ARCHIVE
TAPES**

Generate a report of all of the COPY1 Archive backup tapes, to identify the tape volume serial, data set name, and file number for all archive files created before 1994. This might be used as a library pull list, to prepull the tape volumes required for a FDRTSEL COPY run. To execute the copy, omit the SIM operand and supply a TAPEOUT DD statement. Use of SIM to generate the pull list is highly recommended for large FDRTSEL runs unless the tapes are in an automated tape library

```
//SELLIST      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR, DSN=FDRABR.ARCHIVE
//TAPEOUT       DD     DUMMY
//SYSIN         DD     *
                COPY
                SELECT  ARCHIVE, COPY=1, BKDEVTYP=TAPE, PDATE=94000
                CONTROL SIM
/*
```

60.28 CONTINUED

**MOVE
ARCHIVES TO
NEW DISKS**

Select the archive backups on 3380 disk volume ARCH01 and move them to a new pool of two 3390 disk volumes ABR100 and ABR200. This job will allocate the files on the new volsers with the appropriate size, copy the files from ARCH01, update the Archive Control File appropriately, and scratch the input files on ARCH01 .

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DISP=SHR , DSN=FDRA BR . ARCHIVE
//TAPEOUT       DD     UNIT=3390 , VOL=SER=( ABR100 , ABR200 ) , DISP=OLD
//SYSIN         DD     *
COPY
  SELECT        ARCHIVE , COPY=1 , BKVOL=( ARCH01 )
/*
```

**REDUCE THE
SIZE OF DISK
ARCHIVES**

Select the archive backups on disk older than 30 days and copy them to new files in the same disk pool. Use the ARCEDIT feature to drop all data sets that are marked as RESTORED; this will also drop all archived data sets which are no longer in the Archive Control File. It may result in considerably smaller archive files on disk.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPEOUT       DD     UNIT=3390 , VOL=SER=( ARC001 , ARC002 ) , DISP=OLD
//SYSIN         DD     *
COPY           DYNARC
  SELECT        ARCDIT , COPY=1 , RESTORED=NO , PDAYS=30
/*
```

**COPY
APPLICATION
BACKUP**

A backup tape created with Application Backup (DUMP TYPE=APPL) can be copied with FDRTSEL as long as the Control File in which the backup is recorded is still on disk. Application Backup, described in [Section 51](#), usually involves putting a DSF backup of the Control File which describes the backups on the tape, at the end of the tape, but the most recent copy of the Control File is usually still on disk. FDRTSEL can update the disk Control File and put a fresh backup of it at the end of the output tape created. In this example, a COPY 2 backup is created from all of the COPY 1 payroll backups recorded in the Control File pointed to by the ARCHIVE DD statement. A DSF backup of the Control File will be created as the last file on the tape using GDG name "PAYROLL.CONTROL.BACKUP2"

```
//COPYSEL      EXEC   PGM=FDRTSEL
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//ARCHIVE       DD     DSN=PAYROLL . ARCHIVE . DAILY(0) , DISP=OLD
//TAPEOUT       DD     DSN=PAYROLL . APPL . BACKUP2 , DISP=( NEW , KEEP ) ,
//              UNIT=3490 , LABEL=( 1 , SL )
//SYSIN         DD     *
  SELECT ARCHIVE , COPY=1 , BKDEVTYP=TAPE
  COPY COPYDSN=YES , ABRCOPY=FLIP ,
        ABRINDEX=PAYROLL
  CONTROL ARCBACKUP=DSF , ARCB1DSN=PAYROLL . CONTROL . BACKUP2(+1)
/*
```

60.29 FDRTSEL BACKUP EXAMPLES

**CONSOLIDATE
BACKUPS**

Although ABR normally stacks backups as multiple files on tape, backups may be spread across tape volumes depending on the setup of your backups. This example shows how to make a COPY 2 of all COPY 1 ABR backups for disk volumes beginning with PROD, SYS, and IPLRES, stacking the backups on as few cartridges as possible. Software compression is used even if the original backups are not compressed (COMPRESS=ALL). By default, all backups in the current generation of the selected volumes will be included in the copy (most recent full-volume backup and all succeeding incremental backups).

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=2048K
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPEOUT       DD     DISP=( , KEEP ) , DSN=DUMMY1 , UNIT=3480 ,
//              LABEL=EXPDT=99000 , VOL=( , , 255 )
//SYSIN         DD     *
COPY           COMPRESS=ALL
SELECT         CATLG , SOURCE=1 , COPY=2 ,
              VOL=( PROD* , SYS* , IPLRES )
/*
```

**INCREMENTAL
BACKUPS
DISK-TO-TAPE**

Some installations do daily ABR incremental backups to disk, perhaps because operators are not available to mount tapes during the night. Later, FDRTSEL can be used to move those backups to tape. This example selects all incremental backups (TYPE=ABR) in the current generation for all DASD volumes whose COPY 1 backup is on disk (BKDEVTYP=DISK), and creates a COPY 2 on TAPEOUT and a replacement COPY 1 on TAPE2OUT (ABRCOPY2=1), scratching the backups on disk.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=2048K
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//TAPEOUT       DD     DISP=( , KEEP ) , DSN=DUMMY1 , UNIT=( CART , 2 ) ,
//              LABEL=RETPD=14 , VOL=( , , 255 )
//TAPE2OUT      DD     DISP=( , KEEP ) , DSN=DUMMY2 , UNIT=( CART , 2 ) ,
//              LABEL=RETPD=14 , VOL=( , , 255 )
//SYSIN         DD     *
MOVE           ABRCOPY2=1
SELECT         CATLG , TYPE=ABR , BKDEVTYP=DISK
/*
```

60.29 CONTINUED

CREATE COPY 2 If you create only COPY 1 backups with ABR (to reduce your run time or your tape drive requirements), you can use FDRTSEL to create the second backup asynchronously once the backups complete. IFNOCOPY=2 directs FDRTSEL to copy only those COPY 1 backups for which no COPY 2 already exists. CYCLES=14 indicates that the most recent 14 backups for every DASD volume should be examined for missing COPY 2 backups; this insures that even if the FDRTSEL job is not run or does not complete for a number of days, it will still create all the required COPY 2 backups. The COPY 2 backups will receive the same expiration dates as the equivalent COPY 1s.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPEOUT       DD     DISP=( , KEEP ) , DSN=DUMMY1 , UNIT=( 3480 , 2 ) ,
//              VOL=( , , 255 )
//SYSIN         DD     *
COPY
  SELECT        CATLG , SOURCE=1 , IFNOCOPY=2 , CYCLES=14
/*
```

**STACK
BACKUP
GENERATION**

At some installations, volumes with little or low activity do not receive weekly full-volume backups, substituting daily incrementals over a long period. However, a full-volume restore of a disk with many incrementals may require excessive tape mounts. To reduce recovery time, you can use FDRTSEL periodically to combine these incremental backups to one tape. This example will copy all the COPY 1 backups in the current generation of volumes starting with LIB and create replacement COPY 1 backups with all backups for a given DASD volume stacked on one tape set. Recovery time is greatly enhanced because the files are stacked on the tape in the order in which they are required for a full-volume restore. Expiration dates of the original backups will be copied to the output tapes.

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSUDUMP      DD     SYSOUT=*
//SYSPRINT      DD     SYSOUT=*
//TAPEOUT       DD     DISP=( , KEEP ) , DSN=DUMMY1 , UNIT=( 3480 , 2 ) ,
//              VOL=( , , 255 )
//SYSIN         DD     *
COPY
  SELECT        CATLG , SOURCE=1 , COPY=1 , VOL=LIB*
  CONTROL       STACK=VOLUME
/*
```

**DISASTER
RECOVERY
COPY 1**

Select all backups in the current generation for all DASD volumes and create a COPY 3 to be shipped off site for disaster recovery. For enhanced restore speed, all backups (full and incremental) for a given DASD volume are stacked on a separate tape set (STACK=VOLUME). FDRTSEL automatically writes the backups on the tapes in the order in which they are needed for a restore. Note that this may require repeated mounting of the input tapes and may take considerable time to complete (see next example).

```
//FDRTSEL      EXEC   PGM=FDRTSEL , REGION=4M
//SYSPRINT      DD     SYSOUT=*
//SYSUDUMP      DD     SYSOUT=*
//TAPEOUT       DD     DISP=( , KEEP ) , DSN=DUMMY1 , UNIT=( 3490 , 2 ) ,
//              LABEL=RETPD=35 , VOL=( , , 255 )
//SYSIN         DD     *
COPY
  SELECT        CATLG , SOURCE=1 , COPY=3
  CONTROL       STACK=VOLUME
/*
```

60.29 CONTINUED

**DISASTER
RECOVERY
COPY II**

If you have a large number of backups to copy, you can reduce the elapsed time by breaking the copy into multiple jobs which run concurrently. Each job copies a set of DASD volumes by volser prefix. In this example, COPY 2 is used as the input to create the COPY 3 for offsite storage (perhaps because it is stored in a robotic tape library). UNIT=2 requests that 2 tape drives be used for mounting input tapes, which will reduce elapsed time when copying full-volume backups.

```
//TSEL001      JOB      .....
//FDRTSEL      EXEC    PGM=FDRTSEL,REGION=4M
//SYSUDUMP     DD      SYSOUT=*
//SYSPRINT     DD      SYSOUT=*
//TAPEOUT      DD      DISP=(,KEEP),DSN=DUMMY1,UNIT=(3490,2),
//              LABEL=EXPDT=99000,VOL=(,,,255)
//SYSIN        DD      *
COPY
  SELECT      CATLG,SOURCE=2,COPY=3,VOL=IMS*
  CONTROL     STACK=VOLUME,UNIT=2
/*
//TSEL002      JOB      .....
//FDRTSEL      EXEC    PGM=FDRTSEL,REGION=4M
//SYSUDUMP     DD      SYSOUT=*
//SYSPRINT     DD      SYSOUT=*
//TAPEOUT      DD      DISP=(,KEEP),DSN=DUMMY2,UNIT=(3480,2),
//              LABEL=EXPDT=99000,VOL=(,,,255)
//SYSIN        DD      *
COPY
  SELECT      CATLG,SOURCE=2,COPY=3,VOL=(CICS*,ONL*)
  CONTROL     STACK=VOLUME,UNIT=2
/*
```


60.30 FDRCOPY/FDRTSEL ISPF INTERFACE

FDRTSEL provides an ISPF dialog panel interface. This interface can be used to generate JCL job steps to be executed in the background, or to execute FDRTSEL simulations in the Foreground.

From the FDR main menu, enter T to get to the FDRTSEL menu.

The primary menu can be used to navigate to any one of the major functions available to FDRTSEL. All sub panels can generate JCL, which the user can copy and save to any library. Additionally, the generated JCL can be submitted directly from the panels, or the job can be edited to change any options directly.

FDRTSEL
Primary Menu
(Panel A.T)

```
----- FDRTSEL/FDRTCOPY Primary MENU -----
OPTION ==>

      A. ARCHIVE   Select ARCHIVE/APPL Backups for COPYing
      C. CATLG     Select ABR Full volume or Incremental Dumps

      D. REBUILD   Make a new ARCHIVE tape from its other copy.
      E. RECLAIM   Reorganize ARCHIVE Backup Data Sets on DISK.

      F. MOVE      Migrate ARCHIVE backups from DISK to TAPE.
      G. RECYCLE   Consolidate ARCHIVE backups on TAPE.

      J. SETUP JCL  FDRTSEL/FDRTCOPY defaults for user

SIMULATE ==>          <Yes|NO>          Generate SIMULATION Only
```

FDRTSEL
ARCHIVE
MENU

The ARCHIVE menu can be used to generate FDRTSEL SELECT ARCHIVE or SELECT ARCDIT functions. Through this menu, you can generate any type of copying of Archive backups or APPLication backups.

```
----- FDRTSEL Select Archive Function -----
OPTION ==>

      S Submit      E Edit      F Foreground Simulate      J JCL Setup
-----
      ARCHIVE / ARCDIT Options -----
COPY ==> 1          <1|2>          Input Copy number to Select From
IF ONLY ==> NO      <Yes|No>       Copy input IF it is the ONLY copy

SDATE ==>          <YYddd|YYYYddd> Files created After this date
SDAYS ==>          <nnnn>         Files created Within last nnn days

PDATE ==>          <YYddd|YYYYddd> Files created Prior to this date
PDAYS ==>          <nnnn>         Files created Before last nnn days

ADATE ==>          <YYddd|YYYYddd> Files created on this date

XDAYs ==>          <nnnn>         Files Expiring within next nnnn days
VOL ==>            Disk Volume(s) of data sets in backup

BKVOL ==>          VOLser(s) on which backup resides
BKDEVTYP ==> TAPE   <ANY|DISK|TAPE> Device type of backup files to select

EXPIRED ==>        <Yes|No>       Files Expired or Not yet Expired
RESTORED ==>       <Yes|No>       Include RESTORED data sets or not.
ARCDIT ==>         <Yes|No>       Perform Backup File Edit function.

ARC/APPL ==> 'ABR.ARCHIVE.TEST2000' Control File
CONTROL ==> SIM Control ParmS
```

60.30 CONTINUED

**FDRTSEL
CATLG MENU**

This Menu is the Primary interface for generating SELECT CATLG functions for FDRTSEL. This menu would be used to generate a JOB to copy Incremental or Full volume Backups via their catlaog entry in the ABR catalog. All options for Select CATLG are available through this menu.

```

----- FDRTSEL Select CATALOG Function -----
OPTION ==>

      S  Submit      E  Edit      F  Foreground Simulate      J  JCL Setup
-----

----- Select CATLG Options -----

SOURCE ==> 1          <1-9>          Input Copy number to Select From
COPY  ==> 1          <1-9>          Output Copy number to assign
IFNOCOPY ==>         <1-9>          Create a copy n if None exists
MAXCYC ==>           <nn >          Maximum cycle number
CYCLES ==> 14        <nn >          Maximum number of cycles
MAXGEN ==> 4         <nnnn>         Maximum Generation number
GEN  ==>             <nnnn>         Specific Generation number
TYPE ==> ABR         <FDR|ABR|BOTH> Select only FDR backups or Only ABR
VOL  ==> TSOWK1      Disk Volume(s) of data sets in backup

```

**FDRTSEL
REBUILD MENU**

This menu can be used to generate a job to Rebuild a damaged ARCHIVE or APPLication Backup tape. Usually you know the volser of the damaged tape, but don't always know where all the copy2 versions are located. FDRTSEL uses the damaged volser reference and finds all of the "other" copies of the files on the damaged tape and creates a new version of the files on the damaged tape.

```

----- FDRTSEL Rebuild Damaged Tape -----
OPTION ==>

      S  Submit      E  Edit      F  Foreground Simulate      J  JCL Setup
-----

----- ARCHIVE Backup Tape Rebuild -----

COPY ==> 1          <112>          Enter the copy number of the backup files
                                   on the damaged tape. (B1 or B2 files)

BKVOL ==> IDPBK0      Enter the TAPE volser that you wish to rebuild
                                   from the other copy of the corresponding files.

ARC/APPL ==> 'TSS.ARCHIVE.TSEL2000'      Control File
CONTROL  ==> SIM      Control Params

```

**FDRTSEL
RECLAIM
SPACE MENU**

This menu can be used to generate a JOB to reclaim DASD space within Archive/APPLication backups on DASD if some of the data sets have been restored or reorganized out of the Control File. FDRTSEL will use the SELECT ARCDIT function to copy only the remaining required tracks from the backup file to a new backup file on DASD and release the unused space. This feature can also be used to move disk backups from one set of volumes to another.

60.30 CONTINUED

```

----- FDRTSEL Reorganize Disk Backups -----
OPTION ==>

      S  Submit      E  Edit      F  Foreground Simulate      J  JCL Setup
----- RECLAIM space from Backups on Disk -----

      COPY ==> 2      <112>      Enter the copy number of the Disk backups
                                  to reorganize.

      ARCEDIT ==> NO      <Yes|No>      Do you wish to reclaim unneeded space in the
                                  backup file on disk?

      RESTORED ==>      <Yes|No>      Retain backups of data sets which have been
                                  Restored (YES) or those not yet Restored (NO)

      DISKPOOL ==> SCR083      <==
                                  Enter the DASD volume serial numbers of the
                                  volumes you wish to Reorganize from the
                                  Archive pool disk

----- End of Panel -----

```

**FDRTSEL
MOVE MENU**

This menu can be used to generate an FDRTSEL job to Migrate Archive or Application backups from disk to tape and scratch the input file on disk.

```

----- FDRTSEL Stage Backups from Disk to Tape -----
OPTION ==>

      S  Submit      E  Edit      F  Foreground Simulate      J  JCL Setup
----- MOVE Backup files off Disk -----

      COPY ==> 2      <112>      Copy number of backups to stage to tape.

      ARCEDIT ==> NO      <Yes|No>      Purge unneeded data sets from each backup
                                  file as it is copied to tape.

      RESTORED ==>      <Yes|No>      Keep only data sets not yet Restored. (No)
                                  Keep only data sets Restored. (Yes)

      - - - - - Criteria for Selection: - - - - -

      IF ONLY ==> NO      <Yes|No>      Copy input IF it is the only copy
      ADATE ==>      <YYddd>      Files created on this date
      SDATE ==>      <YYddd>      Files created After this date
      SDAYS ==>      <nnnn>      Files created Within last nnn days
      PDATE ==>      <YYddd>      Files created Prior to this date
      PDAYS ==>      <nnnn>      Files created Before last nnn days
      EXPIRE ==>      <Yes|No>      Files Expired or Not yet Expired
      XDAYS ==>      <nnnn>      Files Expiring within next nnnn days
      VOL ==>      Disk Volume(s) of data sets in backup
      BKVOL ==> 001*      VOLser(s) on which backup resides

      ARC/APPL ==> 'TSS.ARCHIVE.TSEL2000'      Control File
      CONTROL ==> SIM      Control Params

```

60.30 CONTINUED

**FDRTSEL
RECYCLE
TAPES MENU**

This menu can be used to generate FDRTSEL Jobs to copy Archive or APPLication backups from TAPE to a new set of TAPES. Since only the files that are currently recorded in the Control File are selected, this consolidates many tapes on to a smaller set of tapes. The recommended method is by using a date range with the SDATE and PDATE or SDAYS and PDAYS parameters to consolidate a group of tapes per job. If the tapes are under Catalog Control then no special updating is necessary to notify the tape management system that the input tapes are now available for reuse.

```

----- FDRTSEL ReCYCLE Backups Tape to Tape -----
OPTION ==>

      S  Submit      E  Edit      F  Foreground Simulate      J  JCL Setup
-----
                        Consolidate Archive Tapes
-----
COPY ==> 2           <112>           Copy number of backups to stage to tape.
ARCEDIT ==>          <Yes|No>        Purge unneeded data sets from each backup
                                     file as it is copied to tape.
RESTORED ==>         <Yes|No>        Keep only data sets not yet restored. (No)
                                     Keep only data sets restored.      (Yes)
-----
                        Selection by Date Range
-----
SDATE ==>            <YYddd>         Files created after this date
SDAYS ==>            <nnnn>          Files created within last nnn days
                                     Ending date
PDATE ==>            <YYddd>         Files created prior to this date
PDAYS ==>            <nnnn>          Files created before last nnn days
-----
                        Other Criteria for Selection
-----
ADATE ==>            <YYddd>         Files created on this date
EXPIRE ==>           <Yes|No>        Files expired or not yet expired
XDAYS ==>            <nnnn>          Files expiring within next nnnn days
VOL ==>              Dasd Volume(s) backups from.
BKVOL ==> 001*       Tape Volume(s) of backup files
-----
ARC/APPL ==> 'TSS.ARCHIVE.TSEL2000'      Control File
CONTROL ==> SIM                          Control Params

```

60.30 CONTINUED

**FDRTSEL JCL
SETUP MENU**

This menu allows you to define a JOB card and JCL options specific to your installation. Additionally you can set some of the defaults used by all the other FDRTSEL panels. This menu will be invoked automatically the first time a user attempts to use any of the other FDRTSEL panels.

```

----- FDRTSEL/FDRTCOPY JCL options -----
OPTION ==>

----- JOB Statement Information -----
==> //JSA JOB (ACCOUNT),'NAME',NOTIFY=JS
==> /*
==> /*
==> /*

FDR Program Library for STEPLIB DD (blank if LINKLIST) :
  Steplib ==>

SYSOUT Class ==> *

----- General Information -----
SIMULATE ==> YES          <YesNo>          -- Generate SIMULATION Only -
ARC/APPL ==> 'TSS.ARCHIVE.TSEL2000'         Control File
CONTROL ==> SIM          Control Params

----- TAPEOUT DD Information -----
TAPEOUT ==> 'JXS.FDRTSEL.TAPEOUT'           Output DSN
UNIT ==> X                               Output Device Type
RETPD ==>                               Retention Period OR Expire Date
VOL SER ==>                               Output Disk Volume Serials (opt)

----- TAPE2OUT DD Information (optional) -----
TAPE2OUT ==>                               Second Copy DSN
UNIT ==>                               Copy2 Output Device Type
RETPD ==>                               Retention Period OR Expire Date

----- TSELCKPT DD Information (Optional) -----
TSELCKPT ==>                               Checkpoint DSN
UNIT ==>                               Checkpoint Device Type
TRACKS ==>                               Allocation amount
DISP ==> NEW                           <NEW|OLD|SHR>

----- Pool Disk Information (Optional) -----
POOLDISK ==>                               DSN for pooldisk
UNIT ==>                               Device type for DASD backups
VOLUMES ==>                               Volumes for Pooldisk

```

70.01 SYSTEM MANAGED STORAGE (SMS)

MVS and OS/390 include a storage management automation function called System Managed Storage (SMS). If implemented, System Managed Storage provides centralized and automatic assignment of new data sets to pools of disk volumes, and the assignment of characteristics to those data sets which will enable them to be automatically managed.

IBM also includes functions for the support of SMS in their DASD management software: DSS (known as DFDS or DFSMSdss) and HSM (known as DFHSM or DFSMSHsm).

The FDR/ABR system includes the SMS functions of DSS and HSM, making those IBM products unnecessary when FDR/ABR is installed.

**SMS
OVERVIEW**

This brief summary of System Managed Storage and definition of SMS terms is provided to make it easier for those not familiar with SMS to understand this section. More details on SMS are provided by many IBM manuals and education courses which discuss and document it.

SMS provides several significant enhancements over normal data management:

- automatic allocation of new data sets on disk volumes selected by centralized rules
- automatic assignment of management characteristics, so that DASD management systems (such as ABR) can make decisions about the management of data sets at the data set level (rather than at the volume or data set group level).

SMS classes are the mechanism by which SMS makes its decisions about allocation and management. The available class names, and their meanings, are assigned by the storage administrator. When creating new data sets, the classes to be associated with it can be requested by the user (through JCL and IDCAMS parameters), or classes can be assigned through defaults provided by RACF, or by routines coded by the storage administrator.

Data Class allows common characteristics to be assigned to data sets without having to be re-specified. For example, all JCL data sets might be assigned to a data class CNTL which specifies RECFM=FB,LRECL=80,DSORG=PO and other characteristics. They are also used to select special data set formats such as striped and compressed. Data Classes are used only during the creation of new data sets.

Management Class defines how the data set is to be managed by the inactive data manager (ABR or HSM). It includes things like backup frequency and days to retain inactive data sets before archiving. Without SMS, these must be specified by control statements in ABR itself, and usually apply to many data sets. The management class allows the inactive data management to be controlled separately by data set.

Storage Class is used to influence the decision by SMS of what disk volume a data set is to be allocated onto, and can be used to request special types of volumes (such as cache storage, or dual copy). If a data set has a storage class assigned, it will be allocated onto a SMS-managed disk volume. If no storage class is assigned, the data set will be allocated by standard MVS techniques onto a non-SMS volume.

Storage Group consists of one or more SMS-managed disk volumes. Every SMS disk volume belongs to exactly one storage group. A specific storage group cannot be requested by the user, but if a data set has a storage class assigned, a SMS routine will assign one or more storage groups. SMS will then select one volume from among the volumes in those storage groups.

CONTINUED . . .

70.01 CONTINUED

**SMS
OVERVIEW
(continued)**

ACS (Automatic Class Selection) Routines are coded by the SMS storage administrator at each installation. There are ACS routines for the data class, management class, storage class, and storage group. In each ACS routine, decisions are made about the class or group to be assigned to each data set based on many factors (including its name, size, and characteristics, the jobname/userid of the creator, TSO or BATCH, new allocation or restore, etc.). If the user requested specific classes, they are input to the appropriate ACS routine, but the routines can override the user's request. When a data set is being restored from backup or copied, the original classes associated with the data set are passed to the ACS routines (unless overridden by the user).

The SMS classes assigned to a data set (but not the storage group) are stored in both the catalog entry for the data set, and in the SYS1.VVDS.Vvolser on the volume involved. The VVDS was originally for ICF VSAM data sets only, containing VVRs (VSAM volume records), but on SMS-managed volumes, it also contains NVRs (Non-VSAM Records).

SMS has been enhanced since its introduction to support various special types of data sets, including:

- PDSE - PDS Extended
- HFS – Hierarchical File System, used with Open Edition/MVS (also called "UNIX Services for OS/390").
- EF – Extended Format data sets, which includes striped sequential, compressed sequential, compressed VSAM KSDSs, and VSAM KSDSs over 4GB in size.

EF data sets may only be allocated on SMS-managed volumes, but IBM has released enhancements to DFSMS 1.4 (OS/390 V2R4) and above which allow PDSE and HFS data sets on both SMS and non-SMS volumes.

70.10 SMS VOLUME OPERATIONS**FULL-VOLUME OPERATIONS**

FDR and ABR full-volume dumps and restores (TYPE=FDR) and COMPAKTOR operations support SMS volumes. During full-volume dumps, SMS class information from VVRs and NVRs is recorded in control records at the beginning of the backup data set (along with DSCBs and other VVDS information).

During full-volume restores and COMPAKTions, checks are made to ensure that volumes that were SMS-managed when they were dumped can only be restored to current SMS-managed volumes on a system with SMS active, and that non-SMS backups can only be restored to non-SMS volumes. Since SMS data sets must be cataloged, SMS volumes cannot be restored to new volume serials, since that would create uncataloged versions of the data sets on the volume, so CPYVOLID=YES is forced for the restore or COMPAKTion of a SMS-managed volume.

However, since the volume being restored may have changed since the backup was taken, some of the data sets on the volume may not be cataloged or may be cataloged to another volume. Neither FDR, ABR, or COMPAKTOR make any attempt to ensure that the data sets being restored are properly cataloged. Since all data sets on a SMS volume must have entries in the VVDS (either VVRs or NVRs), the IDCAMS command DIAGNOSE can compare the VVDS to the catalog and identify any data sets which are improperly cataloged. The IDCAMS command DELETE NOSCRATCH can be used to delete inaccurate catalog entries, and DEFINE RECATALOG may be used to rebuild accurate catalog entries, including the SMS class information, for both VSAM and non-VSAM data sets.

Under special circumstances, it may be necessary to circumvent these rules. For example, you may need to restore SMS volumes under a non-SMS system for disaster recovery. The keyword SMSPROT=NONE on a FDR or ABR RESTORE or COPY statement bypasses the SMS checks described above.

SMS VOLUME CONVERSION

When implementing SMS, there are two ways of converting data sets to SMS management. A volume can be converted in place, or data sets can be moved from non-SMS volumes to SMS-managed volumes.

IBM literature refers to the CONVERTV function of DSS as the mechanism for converting a volume in place to SMS-management. **The FDR system program FDRCONVT performs the same function, making DSS unnecessary.** FDRCONVT can convert a volume by invoking the SMS ACS routines to assign storage and management classes to every data set on a volume. It can also convert a volume back to non-SMS. [See Section 70.30](#) for complete details.

DSS CONVERTV **must not** be used on a volume initialized for ABR, since it will incorrectly modify the ABR model: **use FDRCONVT.**

FDRCOPY MOVE (described in [Section 21](#)) or FDRDSF DUMP and RESTORE ([Section 20](#)) can be used to move data sets from non-SMS volumes to SMS-managed volumes (assigning storage and management classes) or back to non-SMS, using the techniques described below.

Usually, moving data sets to SMS volumes with FDRCOPY is preferable to using FDRCONVT, if spare DASD volumes are available. Converting volumes requires that an entire volume be converted at one time and that any problems which would prevent conversion (such as uncataloged data sets or ineligible data set types) be resolved. If the volumes being converted contain parts of multi-volume data sets, then all of the other volumes on which those data sets reside must also be converted. A FDRCOPY MOVE allows groups of data sets to be converted in a less drastic manner, and allows for an easier fall-back to non-SMS.

70.11 SMS DATA SET OPERATIONS**DATA SET OPERATIONS**

Data set operations in the FDR system include DSF dumps and restores, FDRCOPY copies and moves, ABR incremental backups, ABR archive backups, FDRAPPL application backups, ABR incremental and archive data set restores, and FDRAPPL data set restores.

For all data set dumps, SMS class information from VVRs and NVRs is recorded in control records at the beginning of the backup data set (along with DSCBs and other VVDS information), and ABR operations can optionally use SMS management class attributes for data set selection (see [Section 70.12](#)).

However, data set restores (including the output side of FDRCOPY operations) are significantly changed on a system with SMS active.

On all systems, SMS or not, an output volume will be chosen. Complete rules for each type of restore are documented in [Sections 20](#) (DSF), [21](#) (FDRCOPY), [50](#) (Volume Backup restore) [Section 51](#) (Archive Backup restore) and [52](#) (FDRAPPL restore). In general, the target volume may be the volume to which the data set is currently cataloged, the original volume from which it was dumped, a volume specified by a DD statement or a volume identified by an NVOL= parameter specified by the user (or by the ABR RESTORE ALLOCATION LIST). If the data set is found in the VTOC of the target volume, the restore/copy/move will be done to that pre-allocated data set and SMS will not be involved.

However, if the data set is NOT found in that volume's VTOC, and SMS is active on the system, SMS will be invoked to decide how and where the data set should be allocated.

SMS CLASS ASSIGNMENT

First, the SMS storage class ACS routine will be invoked to determine if the data set should be SMS-managed. A storage class name will be passed in to the ACS routine:

- if the STORCLAS= operand was specified by the user, that value will be passed.
- if the NULLSTORCLAS operand was specified by the user, a null (blank) class name will be passed.
- if the data set being restored or copied/moved was SMS-managed, and neither operand was given, the original storage class of the data set will be passed.
- if the original data set was not SMS-managed, a null class name will be passed.

The storage class ACS routine may honor the class passed in, or may choose to override it with a different class or a null (blank) class name.

If a null class name is returned, the data set will be non-SMS, and FDR will attempt to allocate it on the target volume previously selected. Note that If that target volume is a SMS-managed volume, but no storage class was assigned, the allocation will fail. If the data set is a type that can exist only on SMS-managed volumes, the allocation will also fail.

If a storage class was assigned, the data set will be allocated as SMS-managed. The SMS management class ACS routine will be invoked next. A management class name will be passed in to the ACS routine:

- if the MGMTCLAS= operand was specified by the user, that value will be passed.
- if the NULLMGMTCLAS operand was specified by the user, a null (blank) class name will be passed.
- if the data set being restored or copied/moved was SMS-managed, and neither operand was given, the original management class of the data set will be passed.
- if the original data set was not SMS-managed, a null class name will be passed.

The management class ACS routine may honor the class passed in, or may choose to override it with a different class or a null (blank) class name (which implies that a SMS default management class will be used).

CONTINUED . . .

70.11 CONTINUED

**SMS CLASS
ASSIGNMENT
(continued)**

The SMS data class ACS routine will not be invoked. Data classes are used only during the allocation of totally new data sets. For data sets being restored or copied/moved, all of the data set characteristics implied by its data class have already been determined. However, the data class is retained in the catalog and VVR/NVR for documentation purposes. The data class of a SMS data set created by FDR can be specified by the DATACLAS= or NULLDATACLAS operands, otherwise the original data class will be used during the allocation if the input data set was SMS-managed.

An authorized storage administrator can bypass the invocation of the class ACS routines by specifying the BYPASSACS operand, in order to control directly the classes assigned. In this case, the SMS classes assigned will be those specified by the STORCLAS=, MGMTCLAS=, and/or DATACLAS= operands, or the original classes associated with the data set if not overridden. A storage class *must* be assigned if the data set is to be SMS-managed.

**CONVERSION
TO AND FROM
SMS**

The description above is complex, so you may be unsure what you have to do to convert a data set from non-SMS to SMS, and back again. Actually it is easy:

- the simplest way to convert a data set to SMS is to update your Storage Class ACS routine so that the data set is assigned a storage class. Then simply move the data sets with FDRCOPY, or dump them with FDRDSF, delete the non-SMS version, and restore them with FDRDSF. Since the output data sets have a storage class, they will be allocated as SMS-managed, and SMS will choose an output volume as described below.
- To convert them back to non-SMS, update the Storage Class ACS routine again so that they are assigned a null storage class, then move or dump/restore them again specifying the NVOL= operand to specify on which non-SMS volumes they will be placed.
- If you don't update the ACS routine, then add the STORCLAS= operand on the SELECT statement in a RESTORE or MOVE step. As long as the Storage Class ACS routine doesn't change it back to a null class, it will be allocated as SMS-managed. You may also want to specify the MGMTCLAS= operand.
- If you add the NULLSTORCLAS operand to the SELECT statement, then the data set will be restored as non-SMS as long as the Storage Class ACS routine doesn't assign a storage class. Be sure to specify NVOL=.
- If your ACS routine doesn't let you do what you want, you can bypass it with the BYPASSACS operand on the RESTORE or MOVE statement (if you are authorized to use it).

**SMS VOLUME
ASSIGNMENT**

If a storage class is assigned, the SMS storage group ACS routine will be invoked next, to determine on which SMS storage group(s) this data set should be allocated, using the SMS storage class as input. SMS will build a list of SMS-managed disk volumes in the selected storage group(s) and attempt allocation on each volume in that list until it is successful (or fails on each).

A FDR restore or copy to a "like" device (such as 3390 to 3390) is much more efficient than an "unlike" restore (such as 3380 to 3390). So, when the FDR system invokes SMS allocation, it instructs SMS to consider "like" devices first, if there are any in the group(s) selected. If allocation cannot be done on a "like" device, SMS will be instructed to look for an "unlike" device.

For data sets which have special hardware requirements, such as EF (Extended Format) data sets, FDR instructs SMS to consider only such volumes.

An authorized storage administrator can bypass the invocation of the storage group ACS routine by specifying the BYPASSSMS operand, in order to directly control the volume on which the data set will be placed. FDR will attempt to allocate the data set directly on the output volume it has chosen (or the user has specified). If the output volume is SMS, a storage class must be assigned to the data set.

Once SMS allocates the data set, its contents will be restored, along with appropriate DSCB and VVR/NVR data (See [Section 80](#) for details). However, the SMS classes in the data set on disk will not be disturbed (either the classes assigned in the steps just described, or the classes associated with a pre-allocated SMS data set). If a SMS-managed data set is being restored to a non-SMS volume, there will, of course, be no classes associated with it.

CONTINUED . . .

70.11 CONTINUED

**GUARANTEED
SPACE
STORAGE
CLASS**

An SMS storage class assigned to a data set may have an attribute of "guaranteed space" which has two functions:

- 1) it allows a user to specify the volume serials of the SMS volumes on which that the data set is to be placed;
- 2) even if specific volume serials are not specified, for multi-volume data sets it pre-allocates the primary space on EVERY volume chosen by SMS.

However, when FDR allocates an SMS data set, it usually passes a specific volume serial, the serial of the target output volume chosen by FDR (as described earlier in this section). If the data set has a guaranteed space storage class, SMS will try to honor the volume chosen by FDR, rather than choosing a volume by SMS rules. The storage group ACS routine will still be invoked, but if the volume chosen by FDR is not in one of the assigned storage groups, or if the allocation fails on that volume, SMS will not be able to allocate the data set.

So, even if the guaranteed space data set was originally allocated on non-specific SMS volumes, FDR will always try to allocate it to a specific volume. At best this circumvents SMS management of data set placement, and at worst, the data set cannot be restored.

The circumvention is to specify NVOL=volser on the SELECT statement, where "volser" is a volume serial that does not exist on your system (e.g., NVOL=DUMMY). In this case, FDR will NOT pass a volser to SMS, and SMS will be free to choose an output volume. For an ABR restore, there must not be an ALLOCATE statement that applies to the data set. However, if the storage class ACS routine returns a null class, making this a non-SMS data sets, the allocation will fail since FDR has no target volume on which to allocate the non-SMS data set.

Because of the confusion this can cause, INNOVATION recommends that use of guaranteed space storage classes be limited to data sets that truly need them.

70.12 SMS ABR SUPPORT

ABR SUPPORT ABR can be run against SMS-managed volumes, as well as non-SMS volumes, even in the same ABR step.

ABR supports those attributes of the SMS management class and storage group which apply to the ABR architecture. However, the SMS attributes were designed by IBM to support the architecture of HSM; since ABR has a different architecture, some of those attributes do not apply to ABR and are not used. For example, ABR manages retention of backups on a volume level, based on the number of generations or retention assigned for that volume, so the attributes relating to retention of backups are not used. A complete list of the SMS attributes NOT used by ABR appears later in this section.

ABR VOLUME SELECTION The SMS-managed volumes to be included in an ABR Backup, Archive Backup or Superscratch step can be specified by the normal ABR volume selection (DISKxxxx DD statements, the ONLINE or ONLVOL operands, or MOUNT statements), or all the volumes defined to a SMS storage group can be included using a MOUNT statement with the STORGRP= operand. For example,

```
MOUNT STORGRP=PRODDB
MOUNT STORGRP=TSO
```

will include all SMS volumes defined in those two storage groups. Both SMS and non-SMS volumes may be selected in the same ABR step.

Regardless of how the SMS-managed volume was selected, it must be properly initialized for ABR processing. The FDRABRM utility ([Section 50.40](#)) or the ABR ISPF dialog (option A.I.8; [Section 90.42](#) and [50.46](#)) must be used to place an ABR Model DSCB in the VTOC of the SMS-managed volume. On SMS volumes, the ABR model will be a cataloged data set.

A valid SMS storage class must be assigned to the ABR model DSCB but it will not be used for volume selection since ABR must force the model to the designated volume being initialized. Any valid storage class may be used; management and data classes are not used for ABR model DSCBs. In your Storage Class ACS routine, you might include code such as:

```
IF &DSN='FDRABR.V*'
  THEN DO
    SET &STORCLAS=ABRMODEL (or whatever name you want to assign)
  EXIT
END
```

This will assign that storage class to all ABR Model DSCBs. You can also change the SET to

```
SET &STORCLAS=&STORCLAS
```

which will simply honor whatever storage class the caller specified. This will require that the STORCLAS= operand be specified on the FDRABRM ABRINIT or REMODEL statements. Also see "SMS Considerations" in [Section 50.40](#).

ABR will use the ABR model **and** the SMS storage group definition of a SMS-managed volume to determine eligibility for ABR processing. The storage group definition includes flags which indicate if the volumes in the group are eligible for "Auto Migrate", "Auto Backup" and/or "Auto Dump". The definition also includes system names and dump classes, which are not used by ABR.

For full and incremental Volume Backup processing (DUMP TYPE=FDR/ABR/AUTO), the SMS volume must have an ABR model **and** its storage group must indicate that it is enabled for "Auto Backup" and "Auto Dump". If it is enabled for only "Auto Dump" without "Auto Backup" then the volume will be processed only for explicit full-volume backups (DUMP TYPE=FDR) and will be bypassed for other ABR backup operations.

For Archive Backup (DUMP TYPE=ARC) and Superscratch (DUMP TYPE=SCR) processing, the ABR model must indicate that it is enabled for Archive Backup or Superscratch (as appropriate) **and** the storage group must indicate that it is enabled for "Auto Migrate".

Neither the ABR model nor the SMS options are used for FDRAPPL Application Backups (DUMP TYPE=APPL). No volume eligibility checks are done for application backups.

CONTINUED . . .

70.12 CONTINUED

Since the options in the ABR model essentially duplicate those in the SMS storage group, you can bypass the checks on the storage group and use only the ABR options for eligibility checking by coding the operand SMSCONSTRUCT=NO on the DUMP statement.

**ABR MANAGEMENT CLASS
USAGE**

By default, ABR will not use management class attributes of SMS-managed data sets for data set selection and processing; data sets will be selected by normal ABR options and operands as described in [Sections 50](#) (Volume Backups) and [51](#) Archive Backup and Superscratch). If you desire, you can run your normal ABR jobs against SMS volumes and select data sets by normal ABR rules, without using SMS management classes at all.

If the operand SMSMANAGE=YES is specified on the DUMP statement, then ABR will use management class attributes for the selection of data sets from SMS volumes while still using normal selection on non-SMS volumes. If you plan to process SMS and non-SMS volumes in the same ABR step, you may need to specify ABR selection operands for the non-SMS volumes (they will be ignored for SMS volumes with the exceptions noted in the following text).

During ABR processing of SMS volumes, ABR will read the definitions of all management classes into storage. As each data set is read from the VTOC of the SMS volume, its associated management class name will be determined and that name will be name in the table. Data sets with no management class will be managed according to the default management class associated with the SMS configuration (if there is no default management class, data sets with no management class will not be selected, nor will data sets with invalid (not currently defined) classes).

Normally, ABR will examine all SMS data sets on the selected SMS volumes and make decisions based on their management classes. However, you may wish to limit ABR to processing only data sets with specified management classes in a particular ABR step. For example, you may wish to limit an Archive backup step to selecting only management classes which will retain Archived data sets for a certain period such as 30 days. As long as SMSMANAGE=YES is specified, you can do so by specifying the management class names on the DUMP statement with the MGMTCLAS= operand (one or more class names can be given); data sets that do not have one of the specified management class names will be bypassed, and data sets that do match will go through the selection process described in this section.

Warning: when SMSMANAGE=YES is specified, data sets on SMS-managed volumes are selected based ONLY on their management class attributes. You cannot use EXCLUDE statements to exclude certain data sets and you cannot use SELECT statements to modify or add to the SMS attributes. However, if a data set is not selected by SMS rules, it is possible to select it anyway by use of a SELECT statement.

Note that Superscratch and Archive Backup are still two separate operations within ABR, requiring two ABR steps if both functions are desired. Since some management class attributes relate to scratching of expired data sets, and some to Archiving (migration) of data sets, you may wish to run Superscratch on SMS-volumes before running ARCHIVE on the same volumes to avoid Archiving data sets unnecessarily.

The following sections detail the use of management class attributes during the three types of ABR operations (Backup, SUPERSCRATCH and ARCHIVE Backup).

70.12 CONTINUED

ABR BACKUP

For full-volume and incremental ABR Volume Backups, if SMSMANAGE=YES, the management class is used only to exclude certain data sets from incremental backups. Management class attributes are NOT used to manage retention of data sets in full and incremental backups.

SMS management class attributes will have no effect on ABR full-volume backups (DUMP TYPE=FDR or full-volume dumps forced during DUMP TYPE=ABR/AUTO/DSF). As usual, all data sets will be dumped during a full-volume backup.

For TYPE=DSF backups, the data sets selected by SELECT/EXCLUDE statements will be dumped unless their management class attribute "ADMIN OR USER COMMAND BACKUP" is set to NONE, which effectively protects the data set from data set backups.

For TYPE=ABR or TYPE=AUTO backups, data sets will be selected by the normal ABR rules ([See Section 50](#)) or by SELECT/EXCLUDE statements, but they will not be dumped if the management class attribute "ADMIN OR USER COMMAND BACKUP" is set to NONE or if the attribute "AUTO BACKUP" is set to NO.

Some attributes of the SMS storage group and management class are designed for use by IBM's HSM and have no meaning for the ABR system since its architecture is considerably different. In particular, the attributes relating to retention of backups are not used; retention of ABR backups follow the normal ABR rules documented elsewhere. The attributes are listed below relate to backup processing, but will be ignored when ABR is used to manage SMS volumes.

Management Class attributes

BACKUP FREQUENCY

NUMBER OF BACKUP VERSIONS (DATA SET EXISTS)

NUMBER OF BACKUP VERSIONS (DATA SET DELETED)

RETAIN DAYS ONLY BACKUP VERSION (DATA SET DELETED)

RETAIN DAYS EXTRA BACKUP VERSIONS

Storage Group attributes

MIGRATE SYSTEM NAME

BACKUP SYSTEM NAME

DUMP SYSTEM NAME

DUMP CLASS

GUARANTEED BACKUP FREQUENCY

CONTINUED . . .

70.12 CONTINUED

**ABR SUPER-
SCRATCH**

Superscratch (DUMP TYPE=SCR) scratches data sets which are not needed, for which no backup is required. When SMSMANAGE=YES is specified on the DUMP statement, then on SMS volumes Superscratch will scratch data sets which are "expired" according to management class rules:

Here is a brief summary of the scratch rules explained in detail below: ABR will scratch SMS-managed data sets based on the expiration attributes of each data set's management class, as documented by IBM. In addition, if you specify EXPIRED on the DUMP TYPE=SCR statement, data sets with an explicit expiration date assigned by its creator or by SMS will be scratched if that date is past.

If you specify the EXPIRED operand on the DUMP TYPE=SCR statement, then Superscratch will scratch data sets that have reached their expiration date. Data sets will have an expiration date recorded in their Format 1 DSCB if the creator of the data set specified an expiration date (EXPDT=yyddd) or retention period (RETPD=dddd) or if the associated SMS data class specified an expiration/retention. If the expiration date in the DSCB is non-zero, and is less than or equal to today's date, the data set is considered expired and will be scratched if EXPIRED was specified.

If EXPIRED is not specified, then data sets with non-zero expiration dates will not be scratched (with the exception of rolled-off GDGs as described below).

For SMS data sets which have zero for the DSCB expiration date (not specified), then the management class attributes "EXPIRE AFTER DAYS NON-USAGE" and "EXPIRE AFTER DAYS/DATE" will be used to determine if the data set should be scratched (the EXPIRED operand is not required):

- If "EXPIRE AFTER DAYS NON-USAGE" is set to a value, the data set may be scratched if it is "nnn" days since the date it was last referenced (stored in the Format 1 DSCB). This is equivalent to the ABR operand ADAYS=nnn.
- If "EXPIRE AFTER DAYS/DATE" is set to a days value, it may be scratched if it is "nnn" days since the data set was created (also stored in the Format 1 DSCB). This is equivalent to the ABR operand CRDAYS=nnn.
- If "EXPIRE AFTER DAYS/DATE" is set to a Julian date, it may be scratched if that date is less than or equal to today's date. This has no ABR equivalent.
- If both of these attributes are set to NOLIMIT then data sets with zero expiration dates will not be scratched. If one of them is set to NOLIMIT then only the other test will be done, but if both of them are set to values, then **both** of the tests must be satisfied for the data set to be scratched.

If the data set is a GDG generation that is in ROLLED-OFF status (fallen off the GDG limit but not deleted from disk because the GDG does not have the SCRATCH attribute or the generation was not expired), the management class attribute "ROLLED-OFF GDS ACTION" will control additional tests:

- If "ROLLED-OFF GDS ACTION" is blank or MIGRATE, the above rules will be used to determine if it should be scratched. If it is not scratched, MIGRATE will cause it to be ARCHIVED during a subsequent TYPE=ARC run.
- If "ROLLED-OFF GDS ACTION" is EXPIRE, the generation will be scratched unless its expiration date in the Format 1 DSCB is greater than today's date. No other tests will be done.

On SMS volumes when SMSMANAGE=YES is specified, only the preceding rules will be used to select data sets for scratching; SELECT/EXCLUDE statements and any GENERAL ARCHIVE SELECTION CRITERIA will not be used, and SCREXCL statements (the SCRATCH PROTECT LIST) will not protect SMS-managed data sets. If you need to scratch data sets based on other criteria, or need to use SELECT/EXCLUDE, you will need to make a separate TYPE=SCR run with SMSMANAGE=NO. However, temporary data sets will be selected by a SELECT TEMP statement even if SMSMANAGE=NO.

You can select or exclude volumes for Superscratch processing based on allocation thresholds, as described in the following section on VOLUME THRESHOLD PROCESSING but it is not recommended. Superscratch takes very little time, so there is little reason to bypass volumes simply because their allocation percentage is low.

70.12 CONTINUED

**ABR ARCHIVE
BACKUP**

Archive Backup (DUMP TYPE=ARC) moves inactive data sets to less expensive media such as tape or compressed disk. In SMS and HSM terms, this is called MIGRATION, but it is the same thing.

If SMSMANAGE=NO is specified or defaulted on the DUMP statement, then ABR will treat SMS volumes like non-SMS volumes. Management class attributes will not apply, and all of the normal ABR Archive selection detailed in [Section 51](#) will apply.

If SMSMANAGE=YES is specified, then during TYPE=ARC processing of SMS volumes ABR will check the "COMMAND OR AUTO MIGRATE" attribute of the management class of each data set:

- If it is set to BOTH then the last reference date in the Format 1 DSCB of the data set is tested to see if the data set has been OPENED within the interval specified by "PRIMARY DAYS NON-USAGE". If not, it will be archived; if so, ABR will test for SMSCOMMAND=YES as described below.
- If it is set to COMMAND, and SMSCOMMAND=YES was specified on the DUMP statement, then the data set will be passed through the normal ABR SELECT/EXCLUDE statements (if any) to see if it should be selected based on normal Archive selection process as detailed in [Section 51](#) (except that if there is no SELECT/EXCLUDE which matches it, it will not be tested against the GENERAL ARCHIVE SELECTION CRITERIA (if any) on the DUMP statement). If SMSCOMMAND=NO (the default), the data set is excluded from archiving.
- If it is set to NONE, the data set is excluded from archiving (similar to the ABR archive protect list).

Note: if the ABR input includes EXCLUDE statements for entire volumes, such as

EXCLUDE ALLDSN,VOL=XYZ123

those volumes will be totally excluded from processing in that ABR step even if SMSCOMMAND=YES is not coded.

So, when Archiving, if SMSCOMMAND=NO is specified or defaulted, data sets will be selected from SMS volumes based only on the management class criteria, while non-SMS volumes will be processed by normal ABR control statement options. Therefore, SMS and non-SMS volumes can be easily processed in one ABR run. If SMSCOMMAND=YES you can have SELECT/EXCLUDE statements which will apply to SMS volumes as well as non-SMS volumes; however, EXCLUDE statements and PROTECT statements (the ARCHIVE PROTECT LIST) will not prevent a data set from being selected from a SMS volume by management class criteria.

If a data set is a rolled-off GDG generation, and "ROLLED-OFF GDS ACTION" is set to MIGRATE, it will be Archived. For all GDG generations, if "# GDG ELEMENTS ON PRIMARY" is not blank, then any generation which is not within the most recent "n" generations will be archived (just like the ABR MAXGDG= operand). Other GDG generations may also be Archived if they meet the tests above. However, GDG generations that are in a management class for which "# GDG ELEMENTS ON PRIMARY" is not blank will only be selected based on management class criteria; even if SMSCOMMAND=YES is specified, they will not be checked against SELECT commands.

CONTINUED . . .

70.12 CONTINUED

**ARCHIVE
EXPIRATION**

As data sets are Archived, they are recorded in the Archive Control File; the location of their COPY1 backup and their COPY2 backup (if created) are recorded, as well as the expiration date of each copy. On non-SMS volumes, the expiration date of COPY1 and COPY2 is set from values specified in the TAPEX and TAPEXX DD statements (or the DUMP statement) and the expiration will be the same for all data sets Archived in the same ABR step. By default, the same is true for SMS data sets. These expirations are referred to as the "normal ABR expirations" in this discussion.

However, if you specify SMSEXPIRE=YES (or SMSEXPIRE=ALL) as well as SMSMANAGE=YES on the DUMP TYPE=ARC statement, then ABR will set the expiration of COPY1 and COPY2 individually for each SMS data set based on attributes of its associated SMS management class.

- SMSEXPIRE=YES applies only to data sets selected by SMS management class criteria; if SMSCOMMAND=YES is specified and some SMS-managed data sets are selected by DELECT commands, those data sets will receive the normal ABR expirations.
- SMSEXPIRE=ALL applies to all data sets selected from a SMS-managed volume, including those selected by SELECT statement.

Most of the following discussion assumes that you are creating a COPY1 on disk and a COPY2 on tape in one ABR step when Archiving data sets from SMS volumes. If this is not true, see the notes later in this subsection.

The intent of the following rules is to set the COPY1 disk expiration of every data set so it can be recalled from disk for the period that the management class specifies it should be on "LEVEL 1", and so that it reaches its final expiration according to other management class attributes.

The COPY1 expiration will be calculated from the attribute "LEVEL 1 DAYS NON-USAGE" plus the Last Reference Date stored in the Format 1 DSCB of the data set. If the calculated expiration date is already past, the expiration is set to today's date. If "LEVEL 1 DAYS NON-USAGE" is set to "NOLIMIT", the COPY1 expiration is set to the normal ABR COPY1 expiration.

If the data set has an explicit expiration date in its Format 1 DSCB (if the creator of the data set specified a expiration date (EXPDT=yyddd) or retention period (RETPD=ddd) or if the associated SMS data class specified an expiration/retention), then that date will be used for the COPY2 expiration.

If it does not have an explicit expiration, the COPY2 expiration will be the higher of:

- the attribute "EXPIRE AFTER DAYS NON-USAGE" plus the Last Reference Date stored in the Format 1 DSCB of the data set.
- the attribute "EXPIRE AFTER DAYS/DATE" (if it specifies a date, that data set is used; if it specifies a days value, that value plus the Creation Date stored in the Format 1 DSCB is used).

If either of these attributes has a value of NOLIMIT, then it is not used. If both are NOLIMIT, then the expiration is set to the normal ABR COPY2 expiration.

There is a minimum COPY2 retention, to avoid the situation where the calculated retention results in a data set expiring almost immediately. It can be specified on the DUMP statement by the SMSMINRET= operand and defaults to 30 days. If the calculated COPY2 expiration is less than the minimum, it is set to the minimum.

If the calculated COPY1 expiration is higher than the COPY2 expiration, it is set to the COPY2 expiration. Also, if COPY1 is on tape, it is set to the COPY2 expiration.

If you are creating a COPY1 on disk and no COPY2, the COPY2 expiration is still stored in the Archive Control File for later use by FDRTSEL as noted below. If COPY1 is on tape, it will receive the COPY2 ("level 2") expiration; the "LEVEL 1" expiration will *only* be used for backups on disk.

70.12 CONTINUED

**ARCHIVE
EXPIRATION
(Continued)**

Even when SMSEXPIRE=YES or ALL is used, the normal ABR expiration dates for COPY1 and COPY2 will still be placed in the labels of the Archive Backup data sets created on disk and tape, so the backup data sets may be retained until those dates are reached.

For this reason, you may wish to divide your SMS Archive Backup into several jobs or steps, using the MGMTCLAS= operand on the DUMP statement to restrict processing to datasets with management classes whose expirations will be within the normal ABR expirations assigned to the COPY1 and COPY2 backups. For example, for a step whose JCL specifies RETPD=10 on COPY1 and RETPD=60 on COPY2, select those management classes whose LEVEL 1 retention is no more than 10, and whose final expiration is no more than 60 days.

Another option would be to use tape management catalog control (usually specified by EXPDT=99000) on the COPY2 on tape; the FDRARCH REORG utility will uncatalog the tape when all files in it have expired, so that it will automatically be released when there are no longer any active data sets stored on it.

The usage of SMSEXPIRE can be very confusing, so here is an example to help clarify it. In this example, the COPY1 backup on disk will be created with a normal ABR retention period of 30 days, and the COPY2 on tape will be scratched by tape management when all files on it are uncataloged (catalog control). However, the expiration dates of COPY1 and COPY2 for the individual data sets will be set in the Archive Control File according to their specified management classes.

SMSEXPIRE=PRT is the same as SMSEXPIRE=YES except that a list of the SMS-managed data sets and the expirations assigned to them will be printed (SMSEXPIRE=PRT is useful with SIM (simulate) to verify correct operation of the management classes). If you use SMSEXPIRE=ALL and want to print the expirations, specify both operands: SMSEXPIRE=ALL,SMSEXPIRE=PRT.

```
//ABRSMS EXEC PGM=FDRABR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSPRIN1 DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//TAPE1 DD DSN=FDRABR.POOLDISK.ABR1,UNIT=DISK,DISP=OLD,
// VOL=SER=(POOL01,POOL02),LABEL=RETPD=30
//TAPE11 DD DSN=ABR11,UNIT=CART,VOL=(, ,255),DISP=(,KEEP),
// LABEL=EXPDT=99000
DUMP TYPE=ARC,SMSMANAGE=YES,SMSEXPIRE=PRT,SMSMINRET=5,
RECALL=YES,DYNARC,DSNENQ=USE,COMPRESS=ALL
MOUNT STORGRP=USER
```

CONTINUED . . .

70.12 CONTINUED

**ARCHIVE
EXPIRATION
(Continued)**

If the job is run on date 98.100, the COPY1 backup data set will expire on 98.130. If the SMS configuration contains management classes with these names and attributes:

DEFAULT (assigned as the default management class in the SMS configuration)

PRIMARY DAYS NON-USAGE=15	EXPIRE AFTER DAYS NON-USAGE=60
LEVEL 1 DAYS NON-USAGE=30	EXPIRE AFTER DAYS/DATE=NOLIMIT

MANAGE01

PRIMARY DAYS NON-USAGE=8	EXPIRE AFTER DAYS NON-USAGE=35
LEVEL 1 DAYS NON-USAGE=25	EXPIRE AFTER DAYS/DATE=50

MANAGE02

PRIMARY DAYS NON-USAGE=5	EXPIRE AFTER DAYS NON-USAGE=30
LEVEL 1 DAYS NON-USAGE=15	EXPIRE AFTER DAYS/DATE=NOLIMIT

then here are some examples of individual data sets that will be selected for archive based on PRIMARY DAYS NON-USAGE and the actual expiration dates that will be recorded for them:

Dsn=A Crdate=98.040 Expdate=none Lastref=98.080 Mgmtclas=none

uses default management class DEFAULT

COPY1 expires 98.110 (Lastref + LEVEL 1 DAYS NON-USAGE)

COPY2 expires 98.140 (Lastref + EXPIRE AFTER DAYS NON-USAGE)

Dsn=B Crdate=98.084 Expdate=none Lastref=98.084 Mgmtclas=MANAGE01

COPY1 expires 98.109 (Lastref + LEVEL 1 DAYS NON-USAGE)

COPY2 expires 98.134 (Crdate + EXPIRE AFTER DAYS/DATE) *

* results in a higher value than (Lastref + EXPIRE AFTER DAYS NON-USAGE)

Dsn=C Crdate=98.084 Expdate=98.140 Lastref=98.090 Mgmtclas=MANAGE01

COPY1 expires 98.115 (Lastref + LEVEL 1 DAYS NON-USAGE)

COPY2 expires 98.140 (Expdate)

Dsn=D Crdate=98.090 Expdate=none Lastref=98.095 Mgmtclas=MANAGE02

COPY1 expires 98.110 (Lastref + LEVEL 1 DAYS NON-USAGE)

COPY2 expires 98.125 (Lastref + EXPIRE AFTER DAYS NON-USAGE)

Dsn=E Crdate=98.090 Expdate=98.107 Lastref=98.093 Mgmtclas=MANAGE02

COPY1 expires 98.108 (Lastref + LEVEL 1 DAYS NON-USAGE) **

COPY2 expires 98.107 (Expdate)

** COPY1 expiration changed to 98.107 because it exceeds COPY2

Dsn=F Crdate=98.080 Expdate=98.102 Lastref=98.086 Mgmtclas=MANAGE02

COPY1 expires 98.101 (Lastref + LEVEL 1 DAYS NON-USAGE)

COPY2 expires 98.105 (SMSMINRET)

70.12 CONTINUED

**ARCHIVE
EXPIRATION
(Continued)**

The FDRTSEL utility ([see Section 60](#)) has several features to help manage archive backups taken with SMSEXPIRE=YES:

- Since the archives within a given backup file will have varying expiration dates, the ARCEDIT function can be used periodically to copy archive tapes and discard the data sets which have expired, reducing the size of the tape library.
- If you create COPY1 on disk and later copy it to tape with FDRTSEL (specifying SMSMANAGE=YES on the COPY/MOVE statement), it will automatically assign the COPY2 expiration (the final "level 2" expiration) calculated above to the new COPY1. This is true even if no COPY2 was created at backup time.

Also, the FDRARCH utility ([see Section 51.50](#)) has been enhanced to support this. If you run a REORGanization of the Archive Control File with the ENABLE=SMSEXPIRE option specified, it will scratch a COPY1 backup on disk if all the entries that point to it are expired, even if the COPY2 entries are not yet expired. This allows REORG to automatically clean up the backup files on disk when they are no longer required, even if the normal ABR expiration for that backup is not yet reached.

**VOLUME
THRESHOLD
PROCESSING**

If the THRESHOLD= operand is specified on the DUMP TYPE=ARC statement, the percentage of space allocated on a SMS volume will be compared to the high-allocation threshold in the SMS storage group associated with the volume (THRESHOLD=HIGH), or the low-allocation threshold (THRESHOLD=LOW) or a user-specified value (THRESHOLD=nn). If the allocation percentage is below the threshold, then Archiving will be bypassed on the volume during this ABR step. This allows you to bypass Archiving on volumes with sufficient free space.

Thresholding is also supported on non-SMS volumes but the high and low thresholds are stored in the ABR model on those volumes. On SMS volumes, the thresholds in the ABR model can be used instead of those in the storage group if the operand SMSTHRESHOLD=NO is specified.

70.13 SMS ACS ROUTINES

**ACS ROUTINE
VARIABLES**

When the SMS ACS (Automatic Class Selection) routines are invoked for allocation of a data set on a system with SMS active, various read-only variables are available to those routines, for use in deciding what classes (if any) to assign to the data set. These variables and their values are fully described in IBM Storage Administration manuals in the chapter on "ACS Language Reference".

The following are brief descriptions of the variables and their values when SMS is invoked by FDRDSF, FDRCOPY, or FDRABR when restoring or copying/moving a data set. Some of the values are set by FDR itself, and some (such as RACF values) are set by SMS before the ACS routine gets control.

ACS variables for ALL data set allocations by FDR

&ACSENVIR – environment in which the allocation is being done:

'RECALL' – for any restore from Archive Backup (including auto-recall)

'RECOVER' – for all other data set restores

'ALLOC' – for FDRCOPY copy/move

'CONVERT' – for FDRCONVT (SMS conversion in-place)

&ANYVOL/&ALLVOL – the target volume chosen by FDR for the data set allocation. For FDRCONVT, the volume(s) the data set currently resides on.

&APPLIC – RACF application identifier

&DEF_DATACLAS – RACF default data class

&DEF_MGMTCLAS – RACF default management class

&DEF_STORCLAS – RACF default storage class

&DSN – data set name being allocated or converted

&DSNTYPE – data set name type ('HFS', 'LIBRARY' (PDSE), 'PDS', null)

&DSORG – data set organization ('PS', 'PO', 'VS', 'DA', null)

&DSOWNER – RACF data set owner

&DSTYPE – data set type ('GDS', 'PERM', 'TEMP', null)

&EXPDT – expiration date (yyyymmdd)

&HLQ – high-level qualifier of the data set name

&LLQ – low-level qualifier of the data set name

&MAXSIZE – maximum size of data set in KB or MB

&NQUAL – number of qualifiers in the data set name

&NVOL – number of volsers (always 1, except for FDRCONVT).

&RECORG – for VSAM, cluster organization ('KS', 'ES', 'RR', 'LS')

&RETPD – retention period

&SIZE – primary allocation size of data set in KB or MB

&UNIT – unit type (e.g., '3390') for the FDR-selected target volume

ACS variables ONLY available for FDRCOPY copy/move

&ACCT_JOB – accounting information from the JOB statement.

&ACCT_STEP – accounting information from the EXEC statement.

&GROUP – RACF group name

&JOB – JOB name, TASK name, or TSO userid

&PGM – job step program name (usually 'FDRCOPY')

&USER – userid of the person running FDRCOPY.

&XMODE – execution mode ('BATCH', 'TSO', 'TASK')

70.20 SMS QUESTIONS AND ANSWERS

The objective of this section is to answer some of the most common questions asked about the FDR products in an SMS system.

Q: HOW DO I RESTORE SMS VOLUMES AT A DISASTER RECOVERY SITE?

A: At a disaster recovery site, you may need to restore SMS-managed volumes on a starter system that does not have SMS active, or which does not include your volumes in its SMS configuration. In this case, you can tell FDR or FDRABR to restore SMS volumes to non-SMS output disks by adding the SMSPROT=NONE keyword to the RESTORE statement; this tells FDR not to make its usual checks when restoring SMS volumes.

Of course, the volumes will be restored with all of their SMS indicators intact, assuming that you will eventually IPL a system with your normal SMS configuration. At that time the restored volumes will be usable. Until the IPL is done, the data on the volumes can be accessed, but no data sets can be allocated or scratched.

Note that you will receive warning messages from FDR indicating that a restore of a SMS volume to non-SMS is being done, but if they contain the text "BYPASSED", they can be ignored.

Q: HOW DO I CONVERT A VOLUME TO BE SMS-MANAGED?

A: Innovation's recommendation is to convert data sets to SMS management by moving them to SMS-managed volumes using FDRCOPY or FDRDSF DUMP/RESTORE, rather than converting the volumes in-place, whenever possible. Moving the data sets allows SMS to position them on the proper SMS volumes according to the SMS classes assigned; data sets not eligible for SMS will be left on their original volumes and can be handled separately.

Program FDRCONVT, documented in [Section 70.30](#), can be used to convert a volume to SMS-managed (and back) in place, without data movement, as long as all data sets currently on the volume are eligible for the conversion. Note that certain data sets (such as ISAM and uncataloged) cannot be converted to SMS, and certain data sets (such as EF (Extended Format)) cannot be converted to non-SMS.

The CONVERTV function of IBM's DSS products can also be used to do such a in-place conversion. However, **CONVERTV cannot be used on a volume that has been initialized for ABR.** CONVERTV will improperly modify the ABR model DSCB ("FDRABR.Vvolser"). FDRCONVT must be used on an ABR volume.

Q: CAN I USE A FULL-VOLUME RESTORE/COPY TO CONVERT SMS VOLUME TO NON-SMS?

A: No, a full-volume restore or copy (via FDR, ABR, or CPK) always restores all of the SMS indicators in the VTOC and VVDS, so it cannot be used to convert a volume. Normally a full-volume operation requires that an SMS-managed volume be restored/copied to a volume currently defined to SMS (and a non-SMS volume only to non-SMS).

Q: HOW DOES SAR HANDLE SMS VOLUMES?

A: SAR (Stand-Alone Restore) is totally unaware of the SMS or non-SMS status of volumes. Since SAR is normally a full-volume dump and restore, it will dump and restore all SMS indicators on the volume; SAR will restore the volume in its original state (SMS or not). When the operating system is reIPLed, SMS volumes will be recognized by SMS if they are currently defined in an SMS storage group. Unlike FDR, SAR does not check the SMS status of the output volume so it will restore to an SMS or non-SMS volume.

CONTINUED . . .

70.20 CONTINUED

Q: ARE THERE CONSIDERATIONS FOR ARCHIVING FROM SMS VOLUMES?

A: The use of the ABR option MIGRAT=YES is a necessity when Archiving data sets from SMS-managed volumes. MIGRAT=YES places a volser of MIGRAT in the catalog entries of Archived data sets instead of the data set's original volume serial number. Since SMS data sets normally have an NVR in the VVDS of their original volume, many MVS functions which operate on the catalog entries of Archived data sets (such as LISTCAT and GDG processing) will get errors if they try to process an Archived data set and discover that the NVR is missing. However, those functions recognize the volser of MIGRAT to mean that the data set is Archived and do not try to access the NVR.

MIGRAT=YES can be specified on the DUMP TYPE=ARC statement of appropriate ABR jobs. However, MIGRAT=YES can also be set as a permanent option in the ABR option table (ABR ISPF panel A.I.4.4) so that it will be automatically used for all Archive Backup jobs. You may want to use MIGRAT=YES for all Archiving, SMS or not.

MIGRAT=YES can and should be used even if DFHSM or DFSMSHsm is installed on your system.

Q: HOW DO I MOVE AN SMS-MANAGED VOLUME WITH FDR?

A: Occasionally you may want to move an entire SMS volume to a new disk, for performance balancing or because of problems with an HDA. This can easily be done with FDR full-volume copy, or full-volume dump and restore, but there are a few considerations for SMS.

First, you may want to be sure that the volume is not in use when you dump or copy it. For a non-SMS volume IBM has never provided any good mechanism to do so, but on SMS volumes you can change the volume to disabled status which will not allow any new jobs to allocate it. This is done by the following console command:

```
V SMS,VOL(volser),DISABLE
```

This will not affect any jobsteps currently using the volume (unless they try to scratch data sets) but will fail any new jobsteps trying to access data sets on the volume. Unfortunately, DISABLE status also prevents FDR from accessing the volume to dump or copy, so you will need to change the status back to ENABLE just before the FDR step starts and back to DISABLE as soon as the dump/copy begins.

If the SMS volume is shared by more than one CPU or LPAR, you should vary the device offline to every system except the one on which the move is to be done.

Now, you can either do a full-volume COPY or a full-volume DUMP and RESTORE to move the volume. In either case, the output volume may be either an SMS-managed volume (with a different volser) or a non-SMS volume. It must, of course, be the same device type with an equal or larger capacity. If it is a larger capacity, you must rebuild the indexed VTOC with ICKDSF BUILDIX.

An example of a full-volume COPY job is:

```
//FDRCOPY      EXEC  PGM=FDR
//SYSPRINT      DD   SYSOUT=*
//SYSPRIN1      DD   SYSOUT=*
//DISK1         DD   UNIT=DISK,VOL=SER=vvvvv1,DISP=OLD <== INPUT VOLUME
//TAPE1         DD   UNIT=DISK,VOL=SER=vvvvv2,DISP=OLD <== OUTPUT VOLUME
//TAPE11        DD   DSN=backup.vvvvv1,UNIT=TAPE, <== TAPE BACKUP
//              VOL=( , , 99),DISP=( ,CATLG) <== RECOMMENDED
COPY TYPE=FDR,CPYVOLID=YES,DSNENQ=HAVE
```

70.20 CONTINUED

If the output volume is currently non-SMS, add ",SMSPROT=NONE" to the COPY statement. After the COPY, the output volume will be renamed to the input volser and placed offline. At this point you issue these console commands:

```
V cuu,OFFLINE <=== original SMS volume address
M cuu,VOL=(SL,vvvvv1) <=== output volume address
```

and the volume will again be usable by SMS (if you DISABLEd the volume to SMS you must now ENABLE it).

An example of a DUMP and RESTORE is:

```
//FDRDUMP EXEC PGM=FDR
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=DISK,VOL=SER=vvvvv1,DISP=OLD <== INPUT VOLUME
//TAPE1 DD DSN=backup.vvvvv1,UNIT=TAPE, <== TAPE BACKUP
// VOL=( , , 99),DISP=( ,CATLG)
DUMP TYPE=FDR,DSNENQ=HAVE
```

... at this point vary the input disk volume offline ...

```
//FDRREST EXEC PGM=FDR
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=DISK,VOL=SER=vvvvv2,DISP=OLD <== OUTPUT VOLUME
//TAPE1 DD DSN=backup.vvvvv1,DISP=OLD
RESTORE TYPE=FDR,CPYVOLID=YES
```

If the output volume is currently non-SMS, add ",SMSPROT=NONE" to the RESTORE statement. After the RESTORE, the output volume will be renamed to the input volser, but SMS will not yet know about it. You must issue these console commands:

```
V cuu,OFFLINE <=== output volume address
M cuu,VOL=(SL,vvvvv1) <=== output volume address
```

and the volume will again be usable by SMS (if you DISABLEd the volume to SMS you must now ENABLE it).

If the original volume was damaged or has hardware problems so that it cannot be dumped or copied, you can do an FDR or ABR full-volume recovery to a new disk volume, and issue the VARY and MOUNT shown above to reactivate the volume. In this case, other steps may be required to put the back-leveled volume in sync with the system catalogs.

CONTINUED . . .

70.20 CONTINUED

Q: HOW DO I IMPLEMENT SMS MANAGEMENT CLASSES IN ABR?

A: ABR supports selection and management of data sets based on SMS management class attributes. [Section 70.12](#) gives much detail on the operation of ABR when management classes are used, but it is a complicated topic. The discussion below attempts to simplify it.

Actually, implementation of management classes in ABR is not difficult. There is a keyword, SMSMANAGE=YES which can be placed on any ABR DUMP statement; it tells ABR to use SMS management classes whenever an SMS volume is being processed, but it is ignored for non-SMS volumes. If you have existing ABR jobstreams which process SMS volumes, you may be able to implement management class selection with minimal modifications.

The description of ABR's processing of management classes in [Section 70.12](#) is detailed, but in simple terms, SMSMANAGE=YES turns on ABR support for all of the attributes of management classes which relate to actual selection of data sets from disk for backup, Superscratch, and Archive (such as "PRIMARY DAYS NON-USAGE" and "ROLLED OFF GDG ACTION"); they will be used by ABR just as they are described in ISMF and in the IBM storage administration manuals.

For ABR Volume backups, SMS management class support can be used ONLY to exclude certain data sets from backup; ABR does not use the attributes which relate to the retention of backups.

In Superscratch and Archive, the management class attributes which relate to expiration and migration of data sets are supported.

Superscratch (DUMP TYPE=SCR) supports the attributes which relate to "expiration" of data sets. Data sets which are considered expired under the management class rules will be scratched from SMS volumes by a Superscratch step which specifies SMSMANAGE=YES (expired data sets which actually have an expiration date in their Format 1 DSCB in the VTOC will be scratched from SMS volumes ONLY if the "EXPIRED" keyword is also specified on the DUMP TYPE=SCR statement).

Archive (DUMP TYPE=ARC) supports the attributes which relate to selection of data sets for "migration" (such as "PRIMARY DAYS NON USAGE").

In an Archive Backup step, you may optionally add the operand SMSEXPIRE=YES or ALL, which enables ABR support to set the COPY1 and COPY2 expiration dates in the Archive Control File separately for every data set selected for Archive by SMSMANAGE=YES processing, based on the management class attributes "LEVEL 1 DAYS NON-USAGE", "EXPIRE AFTER DAYS NON-USAGE" and "EXPIRE AFTER DAYS/DATE". ABR will then consider the COPY1 (Level 1) and COPY2 (Level 2) backups to be expired according to their individual management classes.

So, if you have existing ABR Archive Backup jobs which process SMS volumes, you can add management class support simply by adding SMSMANAGE=YES (and optionally SMSEXPIRE=YES or ALL) to the DUMP statement; ABR will do management class testing on SMS volumes and use normal ABR Archive keywords for selection on non-SMS volumes. If you are not already running Superscratch on your SMS volumes, you will need to add a new ABR step preceding the ARCHIVE step, with DUMP TYPE=SCR,SMSMANAGE=YES to scratch expired data sets (the volumes must also be enabled for Superscratch in the ABR model).

You may want to separate the processing of SMS volumes from non-SMS volumes. The MOUNT STORGRP= statement can be used to select SMS volumes by storage group name (SMSMANAGE=YES is still required on the DUMP statement). When processing SMS volumes you may also wish to limit processing only to data sets with certain management class names; you can do this by specifying the MGMTCLAS= operand on the DUMP statement.

One last note: ABR Archive offers some selection criteria not supported by the management class (such as IFNOTCAT). Of course you can always Archive based on ABR keywords on an SMS volume just by running a step that does NOT specify SMSMANAGE=YES, but you may want to select by management class AND ABR criteria in one step. In that case, specify SMSCOMMAND=YES on the DUMP TYPE=ARC statement (it is not supported for Superscratch); if the management class does not cause a data set to be selected for Archive, ABR will check the ABR control statement criteria as well.

70.20 CONTINUED

Q: WHAT DO I NEED TO DO TO MY ACS ROUTINES FOR FDR?

A: When coding your SMS ACS (Automatic Class Selection) routines you may need to consider a few things relating to FDR; these considerations also apply to IBM's DSS and HSM.

During the initial allocation of a data set, several ACS routine variables can be tested, including the jobname, program name, userid, RACF group and execution mode (BATCH, TSO, or TASK). Your ACS routines may be coded to base class assignments on these variables as well as on other variables such as the data set name, type of data set, and its size.

However, during a restore of a data set that first set of variables will NOT be set; they will be nulls (this is by IBM's rules, not ours). The reason is obviously that during a restore the value of those variables would reflect the job or task doing the restore and not the original owner of the data set, so they would be meaningless during a restore. If the ACS routines test any of those variables (see [Section 70.13](#) for a complete list), they should bypass those tests if the &ACSENVIR variable is "RECALL" (for restore from archive) or "RECOVER" (for any other restore).

It is also important to know that during an FDR data set restore, the original SMS classes of the data set being restored (if it was SMS managed) will be passed as the initial values of the &STORCLAS, &MGMTCLAS, and &DATACLAS variables (those initial values can be overridden by the user using operands on the SELECT statement for other than an auto-recall restore). For a restore, you may wish to code the ACS routine to accept the initial values in most cases.

For FDRCOPY, the &ACSENVIR variable will be set to "ALLOC" and the other ACS variables will be set as if this was a normal allocation. However, you may still need to code special rules if the &PGM is "FDRCOPY" since the variables will reflect the FDRCOPY step, not the original user.

The Data Class ACS routine will never be invoked for a restore or copy/move. The IBM manual GG24-3403 "Writing ACS Routines" is a useful reference.

Q: HOW CAN I CONTROL THE SMS CLASSES AND VOLUMES USED?

A: Sometimes you want to restore, copy or move data sets, but you don't want SMS to override the classes of the original data sets, or you want to specify the output volumes to be used. In this case, you can use the operands BYPASSACS and/or BYPASSSMS on the RESTORE or COPY/MOVE statements.

BYPASSACS does not invoke the ACS class routines at all. The classes assigned to the data set can be specified by the user on the SELECT statement (e.g., MGMTCLAS=) or the original classes of an SMS-managed data set will be used. The output data set must have a storage class assigned (either the original class or specified by STORCLAS=) to be SMS-managed; conversely, if NULLSTORCLAS is specified the data set will be allocated unconditionally as non-SMS.

BYPASSSMS bypasses SMS storage group ACS routine and volume selection, and allocates SMS data sets directly on the output volumes you specify with DISKx DD statements or NVOL= parameters (the volumes MUST be SMS-managed if the data set has a storage class).

Since these options circumvent normal SMS controls, the use of these operands is restricted by a FACILITY class profile in RACF or equivalent security products. The user must be authorized to the resource name.

```
STDADMIN.ADR.RESTORE.BYPASSACS  -- for restores
STDADMIN.ADR.COPY.BYPASSACS     -- for copy/move
```

These names should be restricted to authorized storage administrators.

CONTINUED . . .

70.20 CONTINUED

Q: HOW DO YOU DELETE AN IMPROPERLY CATALOGED SMS DATA SET?

A: By IBM's rules, SMS data sets must always be cataloged, and non-VSAM SMS data sets must have an NVR (non-VSAM record) in the VVDS. Occasionally an interrupted data set creation or other errors may leave an SMS data set only partially created, where the catalog entry, NVR or VVR, and/or DSCB is missing. Such data sets require manual intervention to clean them up since a normal deletion will fail.

To delete the catalog entry for an incomplete data set, issue the IDCAMS command:

```
DELETE name NOSCRATCH
```

The ability to do DELETE NOSCRATCH under SMS may be protected by a security system profile, so it may need to be done by an authorized user.

To delete the NVR and DSCB for an uncataloged SMS data set, issue the IDCAMS command:

```
DELETE name NVR FILE(DD1)
```

where DD1 is a DD statement pointing to the volume involved. DELETE NVR will fail if the data set is actually cataloged. DELETE NVR will still work if only the NVR or only the DSCB exist for the data set.

To delete the VVR and DSCB for uncataloged VSAM components (SMS or not), for EACH component on the volume issue the IDCAMS command:

```
DELETE name VVR FILE(DD1)
```

Q: HOW CAN I PREPARE FOR SMS CONVERSION WITH FDREPORT?

A: FDREPORT can be used to identify data sets that are not supported by SMS such as ISAM, non-ICF VSAM, unmovable and uncataloged data sets. The following example shows how to identify these data sets to ease the conversion to SMS; the report also shows the catalog status of each data set and the volume to which it is cataloged.

```
TITLE LINE=' DATA SETS NOT ELIGIBLE FOR SMS ON TSO VOLUMES '
XSELECT DSORG=(U,IS,AM),DSORG.NE.(EF),VOLG=TSO      UNSUPPORTED DSORG
XSELECT CATALOG=ERR,VOLG=TSO                        CATALOGED TO ANOTHER VOL
XSELECT CATALOG=NO,VOLG=TSO                          UNCATALOGED
REPORT  FIELD=(DEFAULTS,CATALOG,CATVOL)
PRINT
```

NOTE: A simulation with program FDRCONVT can also be used to identify ineligible data sets by volume. An example is shown in [Section 70.35](#).

Q: CAN FDREPORT HELP ME MANAGE MY SMS DATA SETS?

A: FDREPORT can report and summarize based on SMS classes and select based on SMS storage groups so a variety of reports can be generated on SMS-managed data sets and volumes. The example below selects all SMS-managed data sets (data sets that have an SMS storage class assigned) and sorts by storage class name, reporting all data sets which are assigned to each unique storage class, and summarizing based on storage class giving the total occurrences of each class and the total number of tracks of the data sets with that class. You could also report on management class or data class names by changing the field name in the SORT and SUMMARY statements to MGMTCLAS or DATACLAS.

```
TITLE LINE=' ALL SMS DATA SETS BY STORAGE CLASS NAME '
XSELECT STORCLAS.NE.' '                                *SEE NOTE 1
SORT    FIELD=(STORCLAS,DSN),BREAK=(Y,N)              SUMMARIZE ON CLASS NAME
SUMMARY FIELD=(STORCLAS,SIZE)
REPORT  FIELD=(DEFAULTS,SMSCLASS)
PRINT   ONLINE,INFOMSG=NO
```

NOTE 1: STORCLAS.NE. will select all data sets whose storage class is not blank (all SMS data sets have a storage class).

70.20 CONTINUED

Q: CAN I USE ABR TO IMPLEMENT SMS TAPE MOUNT MANAGEMENT?

A: IBM has a feature for SMS called "Tape Mount Management" (TMM). Its function is to divert small data sets which are currently being written to tape (leaving much wasted tape) to SMS-managed disk volumes. IBM recommends that one or more SMS-managed disk volumes be dedicated to this purpose (known as "DASD Buffer" volumes) in a storage group by themselves. This storage group should be processed by the DASD management system at frequent intervals so that all or most of the data sets that reside there can be migrated directly to tape, resulting in efficient tape utilization since many such data sets will occupy one tape.

Tape Mount Management can be implemented with FDRABR. IBM implemented a new attribute in the SMS storage group definition which tells HSM to process the DASD Buffer volumes every hour regardless of other DFHSM options. Since ABR processing is done by batch ABR jobs scheduled by the user, this attribute is not required.

To implement TMM with ABR:

- 1) Define a DASD Buffer storage group to SMS. You will probably want to define it with a large high threshold (such as 90 percent) and a moderate low threshold (such as 50 percent).
- 2) You may need to define several data and management classes to identify and process these data sets. The management classes should specify "PRIMARY DAYS NON USAGE" values of 0 (to allow immediate archiving) or very small values (for data sets which may need to be read back within a day or two).
- 3) Your ACS routines must be modified to recognize tape data sets which are appropriate for TMM and direct them to the TMM classes and storage group. IBM provides a program called "DFP VOLUME MOUNT ANALYZER" which can assist in identifying such data sets.
- 4) A separate ABR job must be coded and scheduled to run against that DASD Buffer storage group at frequent intervals to archive directly to tape (TAPE1 points to a real tape drive). You may want to specify the THRESHOLD=LOW operand in this job so that data sets will be archived from the disk only if the disk is becoming full.

70.30 FDRCONVT OVERVIEW

OVERVIEW FDRCONVT converts a disk volume to or from SMS management without movement of data. It can be used as an alternative to converting data sets to or from SMS management by backing up and restoring them via FDR or FDRDSF or moving them with FDRCOPY.

FDRCONVT can be used to convert a disk volume to SMS or NONSMS status. The volume is first placed in INITIAL status to prevent new data sets from being allocated while conversion is taking place. All of the data sets on the volume are then converted. If any data set cannot be converted, the volume remains in INITIAL status. Once all of the data sets on the volume have been converted, the volume is placed in SMS or NONSMS status, as specified. Conversion of disk volumes to or from SMS status can be simulated prior to actual conversion if desired.

FDRCONVT can place a disk volume in INITIAL status to prevent new data sets from being allocated to the volume in preparation for converting the volume to SMS management. Volumes can be converted to INITIAL status prior to being assigned to an SMS storage group. In order to convert a volume to SMS status, the volume must be assigned to an SMS storage group.

Special processing, such as specifying SMS storage class or management class, can be performed for selected data sets or for all data sets on the volume during conversion to or from SMS status.

**CONSID-
ERATIONS FOR
IN-PLACE
CONVERSION**

FDRCONVT converts data sets to SMS management in-place without moving any data. Conversion in-place is usually faster than converting the data sets by moving the data, especially if the volume contains only a few data sets. However, there are several disadvantages to this method of conversion:

Conversion in-place does not allow the storage management subsystem to determine where the data sets are to be placed. It is a user responsibility to insure that the SMS storage and management classes assigned to the data sets being converted in-place are appropriate. For example, you must insure that a storage class with the AVAILABILITY=CONTINUOUS attribute is not assigned to any data sets on disk volumes being converted to SMS which are not dual-copy devices.

Because data sets are not moved by FDRCONVT, conversion of a volume in-place may result in the high or low ALLOCATION/MIGRATION THRESHOLD values defined in the storage group for the volume to be exceeded. If the high THRESHOLD value is exceeded, the system may not allocate any more data sets to the volume, or data sets could be migrated off the volume to meet the THRESHOLD values.

The VVDS must be able to expand sufficiently to accommodate class information for all data sets being converted. If there are many non-VSAM data sets on a volume, the VVDS may not be able to get enough space or enough secondary extents to convert all of the data sets. In this case, some data sets must be moved off the volume in order to allow the VVDS to expand. Also note that performance may be poor on a volume on which the VVDS is in a large number of extents.

If any data sets on a volume cannot be converted to the desired status (SMS or NONSMS), the volume will remain in INITIAL status. INITIAL status will prevent any new data sets from being allocated on the volume. Also, when converting a volume containing multivolume data sets to NONSMS status, any volumes not specified in the VOL= parameter which contain extents of multivolume data sets will be converted to INITIAL status, since these volumes will contain both SMS and NONSMS data sets.

CONTINUED . . .

70.30 CONTINUED**CONVERSION
BY DATA
MOVEMENT**

The FDRABR, FDRCOPY, and FDRDSF programs can also be used to convert data sets to SMS management by moving data to SMS-managed volumes. The target disk volumes should be initialized with ICKDSF using the INIT command and the STORAGEGROUP and INDEX operands. This will create a VTOC and VTOC index of the desired size and place the volumes in SMS status. The SYS1.VVDS data set will automatically be created when the first VSAM or non-VSAM data set is allocated on a volume, or alternatively, a VVDS can be created explicitly before any data sets are allocated. A large enough VVDS should be defined to accommodate all of the data sets that will reside on a volume.

Once the target volumes are in SMS status, FDRCOPY can be used to move data sets from non-SMS-managed volumes to the SMS managed volumes. FDRDSF and FDRABR can also be used to restore data sets from backups of managed or non-managed disk volumes to SMS-managed volumes.

There are several advantages to converting data sets in this manner. Conversion by data movement allows the storage management subsystem ACS routines to determine the volumes on which the data sets can be placed. This insures that volumes with the proper characteristics are chosen and prevents the converted volumes from having so much data placed on them that the ALLOCATION/MIGRATION THRESHOLDS in the storage group definition are exceeded. The VTOC, VTOC index, and VVDS can be placed adjacent to one another, and a large VVDS can be defined either automatically (the default size for the VVDS is now 10 tracks each of primary and secondary space), or explicitly.

**REQUIRE-
MENTS FOR
CONVERSION**

In order to convert a disk volume to or from SMS status, the volume must be eligible for conversion. Disk volumes which are not assigned to an SMS storage group or which do not have an active indexed VTOC cannot be converted to SMS status. Volumes may be converted to INITIAL status without being assigned to an SMS storage group, but they must have an active indexed VTOC. If a volume is eligible for conversion, it is placed in INITIAL status before conversion of the data sets on the volume is begun.

In addition to the volume being eligible for conversion, all of the data sets which reside on the volume must also be eligible for conversion. If an ineligible data set is found on a volume being converted, that data set is bypassed and conversion continues with the next data set. If all data sets are eligible, the volume is converted to SMS or NONSMS status as specified. If any ineligible data sets were found, the volume remains in INITIAL status. Ineligible data sets must be deleted or moved off the volume before the volume can be successfully converted.

**PREPARING
VOLUMES FOR
CONVERSION
TO SMS**

It may not be practical to attempt to convert all of the data sets on a volume to SMS management at the same time, particularly if new data sets are continually being allocated to the volume. Prior to assigning a volume to an SMS storage group and converting it to SMS status, it may be necessary to quiesce the volume and prevent allocations of new data sets on it. This can be done by using FDRCONVT to place the volume in INITIAL status. New data sets cannot be allocated to a volume in INITIAL status, although existing data sets can be accessed, extended, or deleted. Existing data sets on an INITIAL volume cannot be extended to another volume.

Once the volume is in INITIAL status, no new allocations can take place on it. Data sets that are not eligible for conversion to SMS management can then be deleted or moved to other volumes so that the remaining data sets on the volume can be converted. In this way, the volume can be converted to SMS management over a period of time, so that data sets on the volume remain accessible to users.

As part of the conversion process, FDRCONVT always places volumes in INITIAL status by updating the status indicators in the format-4 DSCB and the VTOC index prior to converting any data sets on the volume to or from SMS management.

70.30 CONTINUED

**CONVERTING
DATA SETS
TO SMS
MANAGE-
MENT**

In order to convert a volume to SMS management, all of the data sets on the volume must be eligible for conversion. The following types of data sets are NOT eligible for conversion to SMS management:

Non-ICF VSAM data sets or catalogs

OS CVOLs and data sets cataloged in OS CVOLs

ISAM data sets

SYS1.STGINDEX data sets

Uncataloged generation data sets (GDGs)

Uncataloged multi-volume data sets

Unmovable data sets (e.g., data sets with DSORG=PSU or POU)

BDAM data sets with OPTCD=A

Active (i.e., ENQueued) data sets

Data sets for which the ACS routines develop a null STORCLAS

Uncataloged data sets are not eligible for conversion to SMS management unless the CATLG=YES parameter is specified, which FDRCONVT to catalog the data sets in the standard order of search catalog.

Multivolume data sets are not eligible for conversion unless either all volumes on which the data sets reside are specified in the VOL= parameter or the MULTIVOL=YES parameter is specified.

A data set is converted to SMS management by calling the ACS routines to determine the storage class and management class that should be assigned to the data set. The ACS routines are called by FDRCONVT in the 'CONVERT' environment (i.e., &ACSENVIR= 'CONVERT'). Data sets for which the ACS storage class routine returns a null SMS storage class are not eligible for conversion to SMS. The data class routines are never called. If the SMS storage class returned from the ACS routines is not null and the data set is otherwise eligible, the data set is converted to SMS management by creating or updating its entry in the VVDS and BCS to add SMS class information and then updating an indicator in the data set's format-1 DSCB to indicate that the data set is SMS managed. Once all of the data sets on the volume have been converted to SMS management, the volume status indicators in the format-4 DSCB and the VTOC index are updated to place the volume in SMS status. Data set aliases will be preserved.

If any data sets on the volume could not be converted, the volume is not converted and remains in INITIAL status. After the ineligible data sets are deleted or moved off the volume, FDRCONVT can be run again to convert the volume to SMS status. When converting to SMS status, data sets that are already SMS managed are not processed again unless REDETERMINE=YES is specified.

70.30 CONTINUED

**CONVERTING
VOLUMES OUT
OF SMS
MANAGE-
MENT**

FDRCONVT can convert volumes in SMS or INITIAL status to NONSMS status. The volume may be may currently be part of a SMS storage group, or not; the deconversion will work either way. If the volume is currently in SMS status, it must have an indexed VTOC. If the volume is in INITIAL status, it must have an indexed VTOC. In addition, all of the data sets on the volume must be either non-SMS managed already or must be eligible for conversion from SMS management.

Certain types of data sets, including EF (Extended Format) files, can reside only on volumes in SMS or INITIAL status. Therefore, volumes containing these data sets cannot be converted to NONSMS status unless they are moved off the volume prior to de-conversion.

Active data sets, i.e., data sets for which a SYSDSN ENQueue fails, cannot be converted out of SMS management.

Eligible data sets are converted out of SMS management by removing the entry in the VVDS for non-VSAM data sets, or by removing the SMS class information from the entry in the VVDS and BCS for VSAM data sets. The ACS routines are not called. Once the VVDS entry for the data set has been updated or deleted, an indicator in the data set's format-1 DSCB is updated to indicate NONSMS status. After all of the data sets on the volume have been converted out of SMS management, the volume state indicators in the format-4 DSCB and the VTOC index are updated to place the volume in NONSMS status.

If any data sets on the volume could not be converted, the volume is not converted and remains in INITIAL status. After the ineligible data sets are deleted or moved off the volume, FDRCONVT can be run again to convert the volume to NONSMS status.

**PROCESSING
SPECIFIC
TYPES OF
DATA SETS****A. UNCATALOGED DATA SETS AND DATA SETS NOT ALIASED**

Multivolume uncataloged data sets and uncataloged generation data sets cannot be converted to SMS management. Single-volume non-VSAM data sets which are uncataloged or are not cataloged in the catalog to which their high level qualifier is aliased cannot be converted to SMS management unless the CATLG=YES parameter is specified.

If the CATLG=YES parameter is specified, uncataloged non-VSAM data sets and non-VSAM data sets cataloged outside the standard order of search are cataloged or recataloged in the appropriate standard order of search catalog (i.e., the catalog to which the data set's high level qualifier(s) is(are) aliased).

The INCAT parameter can be used to specify the catalog to be searched for non-VSAM data sets cataloged outside the standard order of search.

Uncataloged VSAM data sets, e.g., VSAM data sets which have been DELETED with the IDCAMS NOSCRATCH option, cannot be converted to SMS management.

When converting out of SMS management, the only SMS managed data sets that may be uncataloged data sets are GDGs and temporary data sets. Any other data sets which have a format-1 DSCB that indicates SMS-managed but are not catalogued are in error and are not de-converted.

70.30 CONTINUED**B. MULTI-VOLUME DATA SETS**

Multi-volume data sets cannot be converted to or from SMS management unless all of the volumes on which the data sets reside are being converted. You may manually include all the volumes containing the multi-volume data sets in the VOL= parameter, or you may specify the MULTIVOL=YES parameter to cause FDRCONVT to automatically process the multi-volume data sets on all of the volumes on which they reside.

For multi-volume non-VSAM data sets, the first volume on which the data set resides MUST always be specified in the VOL= parameter in order for the data set to be converted to SMS management. For multi-volume VSAM data sets, the FIRST volume on which the base cluster's data or index component or any alternate index cluster's data or index component resides must be specified in the VOL= parameter in order to convert the data set to SMS management. This provision – that the first volume on which the data set or component resides must be in the volume list when converting to SMS management – applies even if MULTIVOL=YES was specified.

All volumes of a multi-volume data set are processed at once when any extent of the data set is encountered. If any volume on which an extent of the data set resides is not eligible for conversion, then the entire multi-volume data set is ineligible for conversion. All volumes which contain extents of a multi-volume data set must be assigned to the same SMS storage group.

For those portions of multi-volume data sets on volumes which were processed because of MULTIVOL=YES (not included in the VOL= parameter), FDRCONVT will flag the Format 1 DSCB of the data set as SMS-managed and will set the status of the volume to INITIAL. To completely convert those volumes to SMS status, you must explicitly run FDRCONVT against them at a later time.

Multi-volume non-VSAM data sets whose catalog entries include candidate volumes are recatalogued with the candidate volumes made non-specific when converting to SMS management (i.e., the catalog entries for the candidate volumes are changed to '*'). When converting multi-volume non-VSAM data sets out of SMS management, the candidate volume entries are removed.

When converting to NONSMS status, the first volume on which the data set resides need not be specified in the VOL= parameter, although MULTIVOL=YES must be specified if any extents of multi-volume data sets reside on volumes not specified via VOL=. If MULTIVOL=YES is not specified, then multi-volume data sets which have extents residing on volumes not specified in the VOL= parameter will still be converted to NONSMS status. However, the volumes which were not specified in the VOL= parameter will be set to INITIAL status, preventing further allocations on those volumes.

C. VSAM DATA SETS

All parts of an ICF VSAM data set, including alternate indexes, are converted at once when any component of the data set is encountered. All parts of a VSAM data set, including alternate indexes and paths, must be cataloged in the same catalog using an alias in order for the data set to be converted.

70.30 CONTINUED**D. GENERATION DATA SETS**

Uncataloged generation data sets cannot be converted to SMS management. Catalogued GDGs can only be converted to SMS management if they are catalogued in an ICF catalog.

Model DSCBs cannot be converted to SMS management because they are uncataloged. Since Data Classes replaced Model DSCBs for SMS-managed GDGs (and can be used even with non-SMS GDGs), the Model DSCBs can usually be deleted.

When converting to NONSMS status, generations that are not active, (i.e., 'deferred' or 'rolled-off' generations) are uncataloged.

E. TEMPORARY DATA SETS

Temporary non-VSAM data sets can be converted to SMS management. The format-1 DSCBs of temporary data sets are marked 'SMS-uncataloged'.

F. THE VTOC, VTOC INDEX, AND VVDS

Flags in the format-4 DSCB and the VTOC index indicate the volume state – SMS-managed, INITIAL status, or NONSMS. The VTOC itself is not a data set and therefore cannot be converted to or from SMS management.

A VTOC index is required on all SMS-managed volumes and volumes in INITIAL status. The VTOC index is converted to or from SMS management along with the other data sets on the volume being converted. When converted to SMS management, the format-1 DSCB for the VTOC index is marked 'SMS-uncataloged'.

A VVDS is required on all SMS managed volumes which are not empty. The VVDS is converted to or from SMS management along with the other data sets on the volume being converted. To conform to IBM conventions, the VVDS will be cataloged into every ICF catalog which is updated for data sets on this volume during the conversion, if it is not already cataloged there.

70.31 PROCESSING OPTIONS AND FEATURES

**SIMULATE
OPTION** FDRCONVT can simulate conversion of volumes and data sets to or from SMS management. This makes it possible to determine from the output of the job what data sets on a volume cannot be converted to SMS management and what SMS storage and management classes will be assigned to those data sets that are eligible for conversion.

**SELECT/
EXCLUDE
CRITERIA** FDRCONVT supports SELECT and EXCLUDE statements that can be used to specify special processing for particular data sets or groups of data sets. Data sets specified in EXCLUDE statements are considered ineligible for conversion. SELECT statements can be used to specify the SMS storage class and management class to be passed as input to the ACS routines for particular data sets. Input storage class and management class can also be specified on the CONVERTVOL or SIMULATE control statements to apply to all data sets processed. Note that the storage class or management class specified may be changed by the ACS routines unless BYPASSACS is also specified.

**PROCESSING
REQUIRE-
MENTS** In order to run FDRCONVT, the storage management subsystem must be active and the volumes which are to be processed must be permanently mounted online and accessible. If volumes are to be converted to SMS management, they must be assigned to an SMS storage group in order for the conversion to be successful.

**STORAGE
REQUIRE-
MENTS** The basic storage requirement for FDRCONVT is 1024K (1M) of virtual storage. This figure includes storage for buffers, control statements, working storage, and program storage. A minimum of 128K of virtual storage must be available below the 16M line. The remainder of the storage may be gotten above 16M. More storage may be required if many ICF VSAM data sets are to be processed. The ICFCORE parameter can be used to increase the storage available for processing ICF VSAM clusters.

**SECURITY
CONSID-
ERATIONS** FDRCONVT checks that the user has RACF FACILITY class authority to the STGADMIN.ADR.CONVERTV resource before beginning conversion of any volumes or data sets. To convert a volume, the user must also have DASDVOL authority to the volume. If the INCAT parameter is used to specify catalogs for data sets cataloged outside the standard order of search, FDRCONVT checks that the user has RACF FACILITY class authority to the STGADMIN.ADR.CONVERTV.INCAT resource. If BYPASSACS is specified, FDRCONVT checks that the user has RACF FACILITY class authority to the STGADMIN.ADR.RESTORE.BYPASSACS resource.

CONTINUED . . .

70.32 FDRCONVT JCL REQUIREMENTS

STEPLIB OR JOBLIB DD STATEMENT If FDR is not in the system linklist, specifies the program library in which FDRCONVT resides. The library must be APF authorized.

EXEC STATEMENT Specifies the program name (PGM=FDRCONVT), region requirement (REGION=1M or more), and optional PARM= operand.

If a PARM field is specified, FDRCONVT will use data specified as the first control statement, which must be a valid CONVERTV or SIMULATE statement. For example,

```
//FDR EXEC PGM=FDRCONVT,PARM='CONVERTV VOL=ABC123'
```

If FDRCONVT is invoked from a user program, Register 1 must follow IBM's convention for passing data from the PARM field.

SYSPRINT DD STATEMENT Specifies the output message data set. Normally a SYSOUT data set. This is a required DD statement.

SYSUDUMP DD STATEMENT Specifies the abend data set. Usually a SYSOUT data set. A SYSUDUMP DD statement should always be included to assist in error diagnosis. If you have the ABEND-AID product from COMPUWARE also include the following so that a fully-formatted dump is produced:

```
//ABNLIGNR DD DUMMY
```

SYSIN DD STATEMENT Specifies the control statements data set. Usually a SYSIN data set. This is a required DD statement unless a control statement is specified in the PARM field of the EXEC statement.

Note: STEPCAT DD statements are prohibited. FDRCONVT will abend if a STEPCAT DD statement is encountered. The INCAT parameter can be used to specify catalogs to be searched for non-VSAM data sets cataloged outside the standard order of search.

CONTINUED . . .

70.33 FDRCONVT CONVERTVOL STATEMENT

CONVERTVOL	STATUS=SMS INITIAL NONSMS	,MGMTCLAS=mgmtclassname
CONVERTV	,BYPASSACS	,MULTIVOL=NO YES
SIMULATE	,CATLG=NO YES	,PRINT=ALL INELIG
SIM	,DSNENQ=USE HAVE TEST NONE	,REDETERMINE=NO YES
	,ENQ=RESERVE ON OFF	,SELTERR=YES NO
	,ICFCORE=nnnnnnnn	,STORCLAS=storclasname
	,INCAT=(catalogname,catalogname,...)	,VOL=(vvvvvv,vvvvvv,...)
	,MAXCARDS=nnnnn	

CONVERTVOL or SIMULATE COMMAND This control statement is required. Multiple CONVERTV or SIMULATE statements are allowed. The statements are processed in the order in which they appear in the control stream.

OPERANDS STATUS=

Specifies the desired status of the volume when conversion or simulation is complete. If the volume is already in the desired status, no action is taken and the disk remains unchanged unless STATUS=SMS and REDETERMINE=YES are specified.

SMS – specifies that the volume is to be converted to SMS format. This is the default. The volume must have an indexed VTOC and must be defined to an SMS storage group in order for conversion to be successful. The volume is first placed in INITIAL status and all data sets on the volume are then converted. If any data set is ineligible for conversion, the volume remains in INITIAL status. If the volume is already in SMS status, no action is taken unless REDETERMINE=YES is also specified.

INITIAL – specifies that the volume is to be placed in INITIAL status. When a volume is in INITIAL status, no new allocations can take place on the volume. INITIAL status is used to prevent new allocations in preparation for conversion to SMS format. Volumes in INITIAL status may contain both SMS and non-SMS data sets. If the volume is already in initial status, no action is taken.

No data sets are converted to or from SMS format when a volume is placed in INITIAL status.

When a volume is converted to INITIAL status from NONSMS status, it must have an indexed VTOC. The volume need not be assigned to an SMS storage group to be converted to INITIAL status, but a storage group must be assigned before the volume can be converted to SMS format. If the volume is specified as part of a group (i.e., the '*' mask character is used in the VOL= parameter) then the volume must be defined to an SMS storage group in order to be converted to INITIAL status from NONSMS status.

NONSMS – specifies that the volume is to be converted to NONSMS status. It does not matter whether the volume is assigned to an SMS storage group unless the volume is specified using the '*' mask character in the VOL= parameter. In this case the volume must NOT be defined to an SMS storage group in order to be de-converted.

CONTINUED . . .

70.33 CONTINUED

BYPASSACS

Specifies that the ACS routines are not to be called when converting data sets to SMS management. If BYPASSACS is specified, STORCLAS must also be specified on the CONVERTVOL or SIMULATE command or the SELECT control statement(s) to provide valid SMS storage classes for the data sets being converted. If a valid SMS storage class is not available for a data set, the data set cannot be converted to SMS management.

The user must have RACF FACILITY class authority to the STGADMIN.ADR.RESTORE.BYPASSACS resource in order to specify BYPASSACS.

CATLG=

NO – specifies that uncataloged data sets or data sets not cataloged in the standard catalog search are not to be cataloged or recataloged when converting a volume to SMS management. These data sets are not eligible for conversion to SMS if CATLG=NO is specified. This is the default.

YES – specifies that uncataloged data sets or data sets not cataloged in the standard catalog search order should be cataloged or recataloged in the appropriate catalog when the volume is converted to SMS management.

Note: If non-VSAM data sets are cataloged outside the standard catalog search order, the INCAT parameter must be specified to indicate what catalogs the data sets can be found in.

The CATLG parameter is ignored unless STATUS=SMS is specified.

Multi-volume uncataloged data sets are never considered eligible for conversion, even if CATLG=YES is specified.

DSNENQ=

Specifies the type of ENQ to be performed on the data sets on the volume. This parameter is optional. Any data set for which the ENQ fails is not eligible for conversion. The default for actual conversion is USE, i.e., the data set is to be enqueued EXCLUSIVE and the operator is not to be prompted. If conversion is only being SIMulated, the default is TEST.

USE – the data set is enqueued. If the ENQ fails the data set is not converted. The operator is not prompted. This is the default.

HAVE – the data set is enqueued, and if the ENQ fails, FDRCONVT will issue a message to the operator (FDRW27). The operator can respond WAIT, NOWAIT, or RETRY. If WAIT is specified, FDRCONVT will wait for the data set to become available. The job could time out. If NOWAIT is specified, FDRCONVT will issue a warning message and the data set will not be converted. If RETRY is specified, FDRCONVT will try the ENQ again.

TEST – The data sets will only be tested to see if they are active. They will not be enqueued. Data sets for which the ENQ fails are considered ineligible for conversion. Note that a user could access the data set after FDRCONVT has issued the ENQ.

NONE – No data set ENQ will be issued.

Caution: when converting a shared DASD volume, the volume should be offline to other systems unless a cross-system ENQ facility is available and the SYSDSN QNAME is broadcast across systems.

70.33 CONTINUED

ENQ=	<p>RESERVE – specifies that FDRCONVT is to issue a RESERVE on the disk volume. This will lock out a SHARED DASD system from accessing this pack, unless a cross-system ENQ facility is available. This is the default.</p> <p>ON – specifies that FDRCONVT is to issue an ENQ for the VTOC.</p> <p>OFF – specifies that FDRCONVT should not issue an ENQ for the VTOC.</p> <p>The ENQ/RESERVE is held only while the F4DSCB and VTOCIX are being updated, rather for the duration of the conversion process. The ENQ/RESERVE may be issued more than once (e.g., when converting the volume from NONSMS to SMS, an ENQ/RESERVE is issued when the INITIAL indicator is set on, and again when the SMS indicator is set on). Because a status of INITIAL prevents allocations on the volume, it is not necessary for FDRCONVT to maintain an ENQ on the VTOC while data sets are being converted.</p>
ICFCORE=	<p>Specifies that FDRCONVT is to increase the size of the tables used to store the ICF VSAM clusters and component names. The value is specified in bytes and must be large enough to contain all of the VSAM names.</p> <p>Note: Specifying ICFCORE= will increase the FDRCONVT memory requirement by the value specified. The default value imposes no additional memory requirement.</p> <p>The default value is 524288, which normally holds about 1100 ICF VSAM components.</p>
INCAT=	<p>Specifies the catalogs to be searched for non-VSAM data sets catalogued outside the standard order of search if CATLG=YES is specified when converting to SMS format.</p> <p>The INCAT parameter is ignored unless STATUS=SMS and CATLG=YES are also specified.</p> <p>In order to specify the INCAT parameter, the user must have RACF facility class authorization to the STGADMIN.ADR.CONVERTV.INCAT resource.</p> <p><i>Caution: the INCAT parameter MUST be specified if there are data sets cataloged outside the standard order of search and CATLG=YES is specified. If INCAT is not specified and there are data sets cataloged outside the standard order of search, they will appear to be uncataloged to FDRCONVT. If CATLG=YES is also specified, then these data sets will be cataloged in both the original catalog and the standard order of search catalog when conversion is complete.</i></p>
MAXCARDS=	<p>Enables FDRCONVT to accept additional SELECT/EXCLUDE statements. Default is a maximum of 250 SELECT or EXCLUDE statements for each CONVERTVOL or SIMULATE statement. An unlimited number of CONVERTVOL or SIMULATE statements may be specified.</p>
MGMTCLAS=	<p>Specifies the SMS management class name which is to be passed to the ACS routines for all data sets unless overridden by a SELECT statement. Must be a valid SMS management class name defined in the active configuration. The default is that the management class name already assigned to the data set is used if there is one. The management class name specified may be changed by the ACS routines and may not be the one assigned to the data set unless BYPASSACS is also specified.</p>

70.33 CONTINUED

- MULTIVOL=** Indicates whether multi-volume data sets are eligible for conversion to or from SMS management if all volumes on which the data set resides are not specified in the VOL= parameter.
- NO** – specifies that multi-volume data sets which have extents that reside on volumes not specified in the VOL= parameter are not eligible for conversion to or from SMS management. This is the default.
- YES** – specifies that multi-volume data sets which have extents that reside on volumes not specified in the VOL= parameter are eligible for conversion to or from SMS management. Volumes not specified in the VOL= parameter will be dynamically allocated if needed.
- When converting to SMS management, the first volume on which a multi-volume data set resides must be specified in the VOL= parameter. For VSAM data sets being converted to SMS management, the FIRST volume on which any of the data or index components of either the base cluster or any alternate index cluster reside must be specified in the VOL= parameter. This restriction does not apply if volumes are being converted to NONSMS status.
- All pieces of a multi-volume data set must be cataloged in the same catalog using an alias.
- All volumes on which a multi-volume data set resides must be assigned to the same SMS storage group and all must be eligible for conversion.
- PRINT=** **ALL** – specifies that information on all data sets processed, including data sets eligible and ineligible for conversion, is to be printed.
- INELIG** – specifies that only information on data sets which are ineligible for conversion to the desired status is printed.
- REDETERMINE=** **NO** – specifies that SMS class information is not to be reset for data sets that have already been converted to SMS management even if the ACS routines return different classes from those already assigned to the data set. This is the default.
- YES** – specifies that all data sets on the volume are to be examined for eligibility for conversion and SMS class information is to be reset for data sets already converted to SMS management if the class information has changed. The REDETERMINE parameter is ignored unless STATUS=SMS is also specified.
- Note:** REDETERMINE=YES can be used on the SIMULATE control statement to insure that the SMS classes assigned to data sets on SMS-managed volumes are valid and are the correct classes that would be assigned by the ACS routines. This function is useful when the ACS routines have been changed or when a volume has been restored at a site where the ACS routines are different from those where the volume was backed up.
- SELTERR=** **YES** – specifies that FDRCONVT is to set a condition code of 12 if a SELECT or EXCLUDE statement is not referenced. This is the default.
- NO** – specifies that FDRCONVT is not to set a condition code of 12 if a SELECT or EXCLUDE statement is not referenced.
- STORCLAS=** Specifies the SMS storage class name which is to be passed to the ACS routines for all data sets unless overridden by a SELECT statement. Must be a valid SMS storage class name defined in the active configuration. The default is that the storage class name already assigned to the data set, if any, is used. The storage class name specified may be changed by the ACS routines and may not be the one assigned to the data set unless BYPASSACS is also specified.

70.33 CONTINUED

VOL=

Specifies the volume serial number of the disk volume(s) to be converted. The volume(s) will be dynamically allocated. Multiple volume serial numbers may be specified in any of the following formats:

- 1.) A list of volume serial numbers may be given, enclosed in parentheses, up to a maximum of 20, e.g., VOL=(TSO001,TSO002,TSO003)
- 2.) A volume group may be specified by placing an asterisk at the end of the volser prefix, e.g., VOL=TSO*
- 3.) The two may be combined, e.g. VOL=(TSO*,PROD*,ABC001)
- 4.) All online volumes may be specified by: VOL=*

The user must have RACF DASDVOL authorization to all of the volumes being converted.

If the '*' masking character is used to specify a group of volumes, the volumes must be assigned to an SMS storage group in order to be converted to INITIAL or SMS status. In order to be converted out of SMS status, the volumes must not be assigned to an SMS storage group if the '*' masking character is used.

70.34 FDRCONVT SELECT STATEMENT

```

SELECT      ALLDSN
S           DSN=datasetname
           DSG=dataset group name

EXCLUDE
X           ,MGMTCLAS=mgmtclassname

           ,STORCLAS=storclassname

           ,VOL=vvvvvv

```

The SELECT or EXCLUDE statement specifies special processing that is to be performed for a particular data set or group of data sets on a volume being converted. SELECT and EXCLUDE statements are processed in the order in which they appear in the control stream.

WARNING: if you include SELECT statements, then only the named data sets will be converted. Since you usually want to convert all the data sets on the volume, you should include a SELECT ALLDSN as the last SELECT, even if it has no parameters, to include all other data sets.

OPERANDS	DSN=	Specifies a unique data set name or cluster to be selected or excluded. From 1 to 44 characters may be specified.
	DSG=	Specifies that a group of data sets matching the characters specified are to be processed. E.g., EXCLUDE DSG=TEST will cause all data sets starting with the characters 'TEST' to be considered ineligible for conversion.
	ALLDSN	Specifies that all data sets on the volume are to be processed. ALLDSN should be specified after all SELECT or EXCLUDE statements for other data sets are specified if all data sets on the volume are to be processed.
	MGMTCLAS=	Specifies the SMS management class name which is to be passed to the ACS routines for the selected data set. Must be a valid SMS management class name defined in the active configuration. The default is that the management class already assigned to the data set, if any, is used. The management class name specified may be changed by the ACS routines and may not be the one assigned to the data set unless BYPASSACS is also specified.
	STORCLAS=	Specifies the SMS storage class name which is to be passed to the ACS routines for the selected data set. Must be a valid SMS storage class name defined in the active configuration. The default is that the storage class name already assigned to the data set, if any, is used. The storage class name specified may be changed by the ACS routines and may not be the one assigned to the data set unless BYPASSACS is also specified.
	VOL=	Specifies the volume(s) to which this SELECT/EXCLUDE statement is to apply. A trailing asterisk can be used as a mask parameter to indicate a group of volumes.

CONTINUED . . .

70.35 FDRCONVT EXAMPLES**EXAMPLE 1 PREPARE A VOLUME FOR LATER CONVERSION TO SMS MANAGEMENT**

Prepare a volume for later conversion to SMS management by using FDRCONVT to place the volume in INITIAL status. Once the volume is in INITIAL status, no new allocations can take place on the volume, and existing data sets on the volume cannot be extended to new volumes. No processing of data sets is done by FDRCONVT when converting volumes to INITIAL status.

```
//CONVT      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
  CONVERTV   VOL=PACK01,STATUS=INITIAL
```

EXAMPLE 2 SIMULATE CONVERSION OF SEVERAL VOLUMES TO SMS MANAGEMENT

Simulate conversion of all of the data sets on several volumes to SMS management. No changes are made to the volumes or the data sets. The output can be used to identify data sets which are ineligible for conversion. The VTOC and data sets are not to be enqueued.

```
//SIMUL      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
  SIM VOL=(PROD*,TEST*),STATUS=SMS,
  ENQ=OFF,DSNENQ=NONE
```

EXAMPLE 3 CONVERT A VOLUME TO SMS MANAGEMENT

Convert a volume containing uncataloged data sets and data sets cataloged outside the standard order of search to SMS management. The uncataloged data sets are to be cataloged. The data sets cataloged outside the standard order of search (i.e., data sets which previously required STEPCATs) are to be recataloged in the standard order of search.

```
//CONVT      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
  CONVERTV VOL=DBPAK1,STATUS=SMS,
           INCAT=(USER.CATALOG.ONE,
                 USER.CATALOG.TWO),
           CATLG=YES
```

EXAMPLE 4 CONVERT MULTI-VOLUME DATA SETS TO SMS MANAGEMENT

Convert volumes containing multi-volume data sets to SMS management. The MULTIVOL=YES parameter indicates that volumes which are not specified in the VOL= parameter can be dynamically allocated. The first extent of each non-VSAM data set and the first extent of each VSAM data or index component or alternate index data or index component must be on a volume specified in the VOL= parameter.

```
//CONVT      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT   DD     SYSOUT=*
//SYSUDUMP   DD     SYSOUT=*
//SYSIN      DD     *
  CONVERTV VOL=(DBPAK1,DBPAK2),STATUS=SMS,
           MULTIVOL=YES
```

CONTINUED . . .

70.35 CONTINUED

EXAMPLE 5 VERIFY THAT SMS CLASSES ASSIGNED TO DATA SETS ARE CORRECT

Simulate conversion of a volume which is already SMS-managed. The REDETERMINE=YES parameter is used to identify those data sets that need to have their SMS class information reset because the ACS routines return classes different from those already assigned.

REDETERMINE=YES may be useful when the ACS routines have been changed, causing new classes to be assigned. If any data sets need to have their class information reset, the SIMULATE command can be changed to CONVERTV to perform the conversion.

```
//CONVT      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      SIMULATE VOL=SMS301,STATUS=SMS,
      REDETERMINE=YES
```

EXAMPLE 6 CONVERT VOLUMES OUT OF SMS MANAGEMENT

All volumes beginning with 'TSO' are to be converted to NONSMS status. Multi-volume data sets can be converted if they reside on volumes specified in the VOL= parameter. Volumes containing EF (extended format) data sets cannot be converted from SMS management.

```
//DECONVT     EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      CONVERTV VOL=TSO*,STATUS=NONSMS
```

EXAMPLE 7 ASSIGN STORAGE AND MANAGEMENT CLASSES TO SELECTED DATA SETS

Convert data sets with high-level qualifiers beginning with 'TEST' or 'PROD' to SMS. All other data sets on the volume are ineligible for conversion. Data sets beginning with 'TEST' are to get a storage class of 'BASE' and a management class of 'MONTHMIG'. Data sets beginning with 'PROD' are to get a management class of DBMIG.

```
//CONVT      EXEC   PGM=FDRCONVT,REGION=1024K
//SYSPRINT    DD     SYSOUT=*
//SYSUDUMP    DD     SYSOUT=*
//SYSIN       DD     *
      CONVERTV VOL=(WRKA*,WRKB*),STATUS=SMS
      SELECT DSG=TEST,STORCLAS=BASE,
      MGMTCLAS=MONTHMIG
      SELECT DSG=PROD,MGMTCLAS=DBMIG
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

80.01 SPECIAL CONSIDERATION OVERVIEW

This section details special considerations, terminology, and rules for programs in the FDR family.

There are two purposes behind this section:

- it centralizes certain explanations when they would otherwise have to be explained in many different sections. For example, the data set name masking and catalog search routine is used in many different parts of FDR, so the documentation in each section refers to this section for details.
- Some supplemental information and more detailed explanations of some topics are here, to make the other sections more usable.

Note: some material which was in the Special Considerations section in previous versions of FDR has been moved to the section to which it applies. For example, the description of naming conventions for ABR backups and for retention of ABR backups has been moved to [Sections 50](#) (Volume Backup), [51](#) (Archiving) and [52](#) (Application Backup).

80.10 DATA SET RESTORE RULES

When a data set is restored by FDRDSF or FDRABR, or copied by FDRCOPY, various pieces of "meta-data" (data which describes the contents of the data set) must also be updated with meta-data from the original data set. This meta-data must be accurate so that the data set is usable. The describing meta-data must be updated even if the output data set is pre-allocated, since it may be different from what was specified by the user during the allocation.

The meta-data always includes the Format 1 DSCB in the VTOC for a data set (for ICF VSAM, there is a DSCB for each component). For ICF VSAM, there is also a VVR (VSAM Volume Record) in the VVDS for each component. For non-VSAM SMS-managed data sets, there is an NVR (Non-VSAM Volume Record) in the VVDS. All of these are updated as appropriate.

In the following text, whenever it refers to "copying a field from the input", the input is either the backup data set for a restore, or the input (original) data set for a copy/move. The indicated fields of the output data set will be updated with the values of those fields in the input data set.

**WHAT IS
UPDATED
IN THE DSCB**

The Format 1 DSCB field names in these tables are taken from the IBM IECSDSL1 macro, used to map DSCBs. Note that some fields, such as DS1SMSFG, will be set when the data set is allocated, either by FDR or pre-allocated by the user, so there is no need to explicitly restore them.

These are fields which are always copied from the input.

DS1VOLSQ	Volume Sequence Number (for multi-volume data sets)
DS1NOBDB	Bytes used in last Directory Block
DS1FLAG1	Special SMS indicators
DS1SYSCD	System Code (Note 2)
DS1DSORG	Data Set Organization
DS1RECFM	Record Format
DS1OPTCD	Option Code
DS1LRECL	Record Length
DS1KEYL	Key Length
DS1RKP	Relative Key Position
DS1TRBAL	Track Balance (unused bytes on the last used track) (Note 4)

The following fields are never modified in the DSCB of the output data set, since the data set may have a different location and/or a different number of extents than the original:

DS1CREDIT	Creation date (Notes 3 and 5)
DS1NOEPV	Number of extents
DS1EXNTS	First 3 Extent descriptors
DS1PTRDS	Pointer to Format 2 DSCB (ISAM) or Format 3 DSCB (more than 3 extents), if any

80.10 CONTINUED

The following DSCB fields are handled specially during the restore or copy/move:

DS1DSSN	Data Set Serial -- this field normally contains the volume serial of the disk on which the data set resides (for multi-volume data sets it points to the <i>first</i> volume of the data set). If DS1DSSN on the input matches the volser of the input disk or the disk that was backed up, then DS1DSSN on the output volume is changed to that output volser. If not, it is restored unchanged
DS1EXPDT	Expiration date -- copied from the input, if zero on the disk (Notes 1 and 5)
DS1REFD	Last Reference Date -- is always set to today's date on RESTORE and COPY, copied from the input on MOVE (Note 5)
DS1BLKL	Block Length - copied from the input unless BLKF= is used to reblock the data set
DS1SMSFG	SMS Indicators (including PDSE and HFS flags) -- if the output data set is SMS-managed, this flag byte is never copied from the input (flags are already set in the output DSCB). When restoring a non-SMS data set to non-SMS, the byte is always copied. When restoring an SMS data set to non-SMS, the byte is copied but the SMS-managed flag is reset.
DS1SCEXT	Extended Secondary Allocation Fields -- they are valid only if a flag is set in the DS1SCALO field. Will be copied from the input only if the extended secondary allocation flag is not set on the output DSCB
DS1DSIND	Data Set Indicators -- copied from the backup except that the update indicator X '02' is always turned on
DS1LSTAR	Last Used Track (relative track number) -- copied from the input except for PDSE and HFS data sets (Note 4)
DS1SCALO	Secondary allocation -- copied from the input, if zero on the disk. If the output data set is preallocated and a secondary allocation quantity was specified, it will not be modified by restore (Note 1)
ABR RESERVED BYTES	Displacement 103 and 104 (decimal) in the DSCB -- these fields are used by ABR Volume Backups; they contain flags and the current cycle number. Depending on the circumstances, <ul style="list-style-type: none"> - the bytes may be zeroed (indicating no current backup) - the bytes may be preserved in the output data set - the bytes may be copied from the input - the bytes may be copied from the input and modified to indicate that the ABR backup read for the restore is the current backup of the data set (if OLDBACKUP is enabled, the old backup information in DS1SYSCD is also modified)

Note 1: If FDR allocates the output data set, DS1EXPDT and DS1SCALO are always copied from the input.

Note 2: If this field is being used to record OLDBACKUP information from ABR Volume Backups, then the restore may do special OLDBACKUP processing on the contents.

Note 3: An FDRCOPY MOVE will copy DS1CREDT from the input data set.

Note 4: For a logical restore (used for unlike disk devices and BLKF=) DS1LSTAR and DS1TRBAL will be recalculated to point to the actual end of the data.

Note 5: options are available to change the processing for the date fields DS1CREDT, DS1REFD and DS1EXPDT. ISPF panel A.1.4.2 contains options for each of these dates to force FDR to always copy the date from the input (the backup file or input data set).

If you require special processing for any of these DSCB fields, contact Innovation for assistance.

80.10 CONTINUED

**FIELDS IN THE
VVDS**

On each volume, the VVDS (VSAM Volume Data Set) contains additional information about ICF VSAM clusters and SMS-managed data sets. The VSAM data is in a record called the VVR (VSAM Volume Record); this includes SMS information for SMS-managed clusters. For SMS-managed non-VSAM data sets, the data is in a record called the NVR (Non-VSAM Volume Record).

The VVR and NVR are each divided into cells and sometimes sub-cells (cells within another cell). Each cell type has a number. A Type 24 subcell is common between VVRs and NVRs, but in general they consist of cells unique to the VVR or NVR.

The text in the following sections describes in general terms what information is restored or modified in each VVDS record. It does not attempt to give details of processing for every field within those cells.

**WHAT IS
UPDATED IN
THE VVR**

A VVR may be a VVDS record type Z or Q; each represents a **component** of a VSAM cluster, so a volume can contain a 2 VVRs for a KSDS with data and index on the same volume. Z records are present on the primary (first) volume of each component of the cluster (the first volume of the data and the first volume of the index, if present). Q records exist on secondary volumes for multi-volume data sets. Q records are also used for imbedded indexes (the IMBED option of IDCAMS DEFINE CLUSTER); in this case a volume might contain 2 VVRs for the index component, a Z or Q for the real index and a Q for the index imbedded in the data component.

VVR header: identifies the record type (Z or Q) and contains:

- Component name
- Cluster name
- Name of the ICF catalog in which the cluster is cataloged
- Base cluster name (same as "cluster name" except for AIXs (Alternate Indexes) where the cluster name is the AIX cluster, while this field points to the base cluster to which the AIX is related.

These will not be changed; if FDR allocated the data set then they are already correct; if not, FDR will not overlay the existing names.

Type 21 cell: exists only in the Z record on the primary (first) volume of each component and contains some descriptive information, including:

- cluster attributes (e.g., SPEED and REUSE)
- sharing attributes (SHAREOPTIONS(n,n))
- space allocation attributes (primary and secondary allocation quantities)
- high-allocated and high-used RBAs (Relative Byte Addresses) for the entire component
- flags indicating if the cluster is an Extended Format (EF) or Extended Attribute (EA) cluster. These can exist only on SMS-managed volumes. VSAM EF clusters can be compressed and can exceed 4GB in size

Most of these fields are copied from the input.

Type 24 subcell: part of the Type 21 cell, it contains SMS information, primarily the SMS storage, management and data class names assigned to the data set. These will not be changed; if FDR allocated the data set then they are already correct; if not, FDR will not overlay the existing classes.

80.10 CONTINUED

**WHAT IS
UPDATED IN
THE VVR
(continued)**

Type 60 cell: exists only in the Z record on the primary volume of each component and contains some basic descriptive information and statistics, including:

- cluster attributes (e.g., INDEXED, IMBED, KEYRANGE, LINEAR)
- free space information
- format information (e.g., CIs per CA, CISIZE)
- usage statistics (e.g., number of CA splits)
- last-updated timestamp
- Number of extents in the entire component

Most of these fields are copied from the input.

Type 23 cell: exists in all Z and Q records on every volume of each component and describes the piece of the component on that volume, including:

- volume/component attributes (e.g., Q record for imbedded index)
- format information (e.g., physical block size, blocks/track)
- high-allocated and high-used RBAs (Relative Byte Addresses) for the component on this volume
- number of extents in the component on this volume
- extent descriptions (starting and ending track addresses) and RBA ranges for each extent

Since this cell describes the cluster as it is allocated on the output volume, most of these fields are not copied from the input.

Type 27 cell: contains information for Extended Format (EF) data sets on SMS-managed volumes. VSAM EF data sets include compressed KSDSs and KSDSs over 4GB in size. The entire cell is restored; it will be created if it doesn't exist.

Type 2A cell: contains information for Extended Attribute (EA) data sets on SMS-managed volumes. This is primarily for data sets used by NFS (Network File System). The entire cell is restored; it will be created if it doesn't exist.

When restoring or copying to a pre-allocated cluster (not allocated by FDR), FDR will verify that certain critical characteristics of the cluster on the output volume match those of the input data set, such as type (INDEXED, IMBED, etc.), allocation (CISIZE, CIs per CA), and attributes (such as EF). Message FDR152 is printed when there is a mismatch; please review that message for all of the attributes which must match.

When a VSAM cluster is restored or moved to an unlike device type (e.g., 3380 to 3390), many of the fields in these cells will be recalculated to reflect the new layout of the cluster on the new device. When VSAM is restored to a unlike disk, it cannot be pre-allocated; FDR must allocate the output cluster, so it controls all of the attributes of the new cluster to meet its needs.

80.10 CONTINUED

**WHAT IS
UPDATED IN
THE NVR**

An NVR (VVDS record type N) exists only on the first or only volume of a SMS-managed non-VSAM data set.

NVR header: identifies the record type (N) and contains the data set name and the name of the ICF catalog in which the cluster is cataloged. These will not be changed; if FDR allocated the data set then they are already correct; if not, FDR will not overlay the existing names.

Type 22 cell: contains some basic indicators. For GDGs, it indicates if the generation is active, deferred, or rolled-out. It also indicates if the data set is a PDSE or HFS data set. There are flags indicating if the data set is an Extended Format (EF) or Extended Attribute (EA) data set. All of these are restored.

Type 24 subcell: part of the Type 22 cell, this contains SMS information, primarily the SMS storage, management and data class names assigned to the data set. These will not be changed; if FDR allocated the data set then they are already correct; if not, FDR will not overlay the existing classes.

Type 28 cell: contains information for Extended Format (EF) data sets. Non-VSAM EF data sets include striped and compressed sequential data sets. The entire cell is restored; it will be created if it doesn't exist.

Type 2B cell: contains information for Extended Attribute (EA) data sets. This is primarily for data sets used by NFS (Network File System). The entire cell is restored; it will be created if it doesn't exist.

80.11 FDR PROCESSING BY TYPE OF DATA SET

This section describes how FDR data set operations (including DSF dump and restore, ABR dump and restore, and FDRCOPY copy and move) handle specific types of data sets.

Since FDRCOPY is essentially a data set dump and restore in the same task, all comments on dump apply to the input side of FDRCOPY and all comments on restore apply to the output side, unless there are specific notes about FDRCOPY.

As described in [Section 80.12](#), a "physical" restore is usually used when restoring a data set unmodified to the same type of disk it was dumped from; tracks are restored exactly as they were on the input disk, except that they may be at new locations on the output disk. A "logical" restore is usually used when restoring to an unlike disk type, or when reblocking of the data set is requested; the data set will be rearranged for the new disk geometry.

Member UNLIKE in the FDR Installation Control Library (ICL) contains the latest information on logical restores; you should review it before moving data to unlike disk.

Unlike some competing products, all FDR backups are "physical" so that exact images of the input data set's tracks are written to the backup. The DSCBs and VVDS records for all data sets dumped are written to the beginning of the backup data set for restore use. An FDR or ABR full-volume backup treats individual data sets the same as a data set backup, except that the default of DATA=ALL instead of DATA=USED is used.

**PHYSICAL
SEQUENTIAL
(PS)**

As the name implies, PS data sets are simple disk data sets which are written and read sequentially (from the start to the end).

DUMP: Only the used tracks in a PS data set (up to the "last block pointer" in the Format 1 DSCB) are dumped, unless the option DATA=ALL is in effect. If the last block pointer is all zeros (which could mean that the data set contains no records, or that it was never opened for output, or that it was updated by an access method that does not maintain the LBP), the entire allocated space will be dumped.

RESTORE: Normally the used tracks of a PS data set will be restored (up to the last block pointer); DATA=ALL will restore all of the allocated tracks and should be used only when restoring from a dump taken with DATA=ALL. A physical restore will restore each track image for each selected track. A logical restore will rearrange and/or reblock the data set and will restore data until a EOF (End-Of-File) record is encountered; DATA=ALL is ignored.

If reblocking is requested (BLKF= operand), and the data set has RECFM=FB or VB the logical records will be reblocked to the new physical blocksize, which must be larger than the original blocksize except when restoring to a disk with a smaller track size than the original disk.

ALLOCATION: if the output data set does not exist on the target output volume, FDR will allocate it. It will be allocated with the same number of tracks as the original unless the CYL=, TRK=, RLSE=, or %FREE= operands are specified (for a logical restore to an unlike device the space may be recalculated). The allocated space may be in fewer or more extents than the original, but the total tracks allocated must equal or exceed the number of tracks to be restored.

MULTI-VOLUME: Since FDR is a volume-oriented product, the pieces of multi-volume PS data sets will be backed up to separate backup files. During the restore, FDR will allocate each piece of the multi-volume data set on a separate volume, so **it must be restored to as many volumes as it was dumped from**. The data set will not be usable until all original pieces have been restored.

**PHYSICAL
SEQUENTIAL
(PS)
(continued)**

EXTENDED FORMAT: On SMS-managed volumes, you can create a new format of PS data set called PSE (Physical Sequential Extended); also called an Extended Format (EF) data set. PSE data sets can be striped (allocated as multi-volume for performance) and/or compressed. FDR can allocate PSE data sets on SMS volumes, and will restore the striping and compression information.

**PARTITIONED
(PO-PDS)**

Partitioned Data Sets (PDS) contain mini-sequential data sets called "members", pointed to by a directory. They are typically used for program libraries, source libraries, and control statement libraries.

DUMP: Only the used tracks in a PO data set (up to the "last block pointer" in the Format 1 DSCB)

80.11 CONTINUED

are dumped, unless the option DATA=ALL is in effect.

RESTORE: Normally the used tracks of a PO data set will be restored (up to the last block pointer); DATA=ALL will restore all of the allocated tracks and should be used only when restoring from a dump taken with DATA=ALL. A physical restore will restore each track image for each selected track; the PDS will not be compressed.

A logical restore will rearrange the data set, and update the PDS directory to point to the new location of each member (note lists will also be updated). Only live members (pointed to by a current directory entry) will be restored, so the PDS will be compressed. Diagnostics will be produced if invalid directory entries or members are found during the logical restore. Logical restore will read the input data only up to the EOF (End-of-File) following the last active member; DATA=ALL is ignored.

If reblocking is requested (BLKF= operand), the new physical blocksize will be stored in the Format 1 DSCB of the PDS, and will be used for new or updated members, but the existing members will not be reblocked during the restore.

PDS directories cannot be expanded in size during a restore. Even if the output data set is pre-allocated with a larger directory, the restore will reset it to its previous size. If you are licensed for FDRREORG, the REORG command of FDRCOPY may be used to compress PDSs and expand directories ([see Section 32](#)).

ALLOCATION: if the output data set does not exist on the target output volume, FDR will allocate it. It will be allocated with the same number of tracks as the original unless the CYL=, TRK=, RLSE=, or %FREE= operands are specified (for a logical restore to an unlike device the space may be recalculated). The allocated space may be in fewer or more extents than the original, but the total must equal or exceed the number of tracks to be restored. PDSs cannot be multi-volume.

Warning: *If a PDS is in multiple extents, the directory must be entirely contained in the first extent; if it isn't the PDS will be unusable. A physical restore does not check whether the first extent of the output data set will hold the PDS directory. A custom zap (C53.0512) is available to fail the restore of a PDS unless FDR can allocate it as a single contiguous extent.*

80.11 CONTINUED

**PARTITIONED
EXTENDED
(PDSE)**

Externally PDSE data sets are used like standard PO data sets, but internally they have 4K blocks and an expandable directory, and reuse space so that they rarely need compression. Originally they were supported only on SMS-managed volumes, but IBM has released enhancements for DFSMS V1R4 (OS/390 V2R4) and above to support them on non-SMS volumes as well.

DUMP: All allocated tracks are dumped.

RESTORE: All allocated tracks are restored. A physical restore restores each track image. A logical restore will rearrange the 4K blocks to fit on the new device; since pointers in a PDSE are by relative block number, not TTR, no other changes are made. If BLKF= is specified, the simulated blocksize used by access methods will be updated in the Format 1 DSCB, but the 4K internal blocks of the PDSE are not changed. An IGWNOTIF macro is issued to purge any PDSE data still in buffers for this PDSE. A restore or copy can be used to convert a SMS-managed PDSE to non-SMS or vice versa.

ALLOCATION: if the output data set does not exist on the target output volume, FDR will allocate it. It will be allocated with the same number of tracks as the original unless the CYL=, TRK=, RLSE=, or %FREE= operands are specified (for a logical restore to an unlike device the space may be recalculated). The allocated space may be in fewer or more extents than the original, but the total must equal or exceed the number of tracks to be restored.

**HIERARCHIAL
FILE SYSTEM
(HFS)**

HFS data sets use the same internal structure as PDSEs but they are used to store UNIX-style files for use with OS/390 UNIX System Services (previously called Open Edition/MVS). Originally they were supported only on SMS-managed volumes, but IBM has released enhancements for DFSMS V1R4 (OS/390 V2R4) and above to support them on non-SMS volumes as well.

DUMP: All allocated tracks are dumped. If HFS=QUIESCE is specified on the DUMP, SNAP, SPLIT or COPY statement, FDR will first try to get a SYSDSN ENQ on the HFS data set. If it is successful, the HFS file cannot be mounted to OS/390 Unix so this is sufficient to insure that no other task can use it during the backup. If the ENQ fails, which probably means that the HFS file system *is* mounted, FDR will attempt to do a "quiesce" function to lock out all users of the file system until the backup is complete. This may cause other OS/390 Unix tasks using that **HFS** file to wait.

FDR jobs which use HFS=QUIESCE must be authorized to use OS/390 Unix services: the userid under which the dump is run must be defined to your security system and it must have a OS/390 Unix UID assigned to it; it does not have to be a Unix superuser. In RACF, the UID is assigned in a OMVS segment attached to the userid or by a default OMVS segment defined by the resource BPX.DEFAULT.USER in class FACILITY; see IBM RACF documentation for details. For other security systems, consult their documentation for the equivalent functions.

If both the ENQ and the quiesce fail, then updates to the HFS file system may occur during the backup, which may make the backup unusable. For data set backups, if ENQERR=BYPASS is specified, the HFS data set will not be backed up if the quiesce fails. If the HFS file system is mounted for read/write on another system, the quiesce may not prevent updates on that system so you should backup such HFS data sets only on a system where they are mounted as read/write.

Quiesce is not used for COMPAKTOR, ABR Archive, or FDRCOPY MOVE. In these functions, HFS data sets are bypassed or made unmovable if the SYSDSN ENQ fails. If you want to move or archive HFS data sets, you must dismount them first.

RESTORE: All allocated tracks are restored. A physical restore restores each track image. A logical restore will rearrange the 4K blocks to fit on the new device; since pointers in a HFS data set are by relative block number, not TTR, no other changes are made. A restore or copy can be used to convert a SMS-managed PDSE to non-SMS or vice versa. If you are restoring to an existing HFS data set, you must unmount the file system before doing the restore.

ALLOCATION: if the output data set does not exist on the target output volume, FDR will allocate it. It will be allocated with the same number of tracks as the original unless the CYL=, TRK=, RLSE=, or %FREE= operands are specified (for a logical restore to an unlike device the space may be recalculated). The allocated space may be in fewer or more extents than the original, but the total must equal or exceed the number of tracks to be restored.

80.11 CONTINUED

See member HFS in the FDR ICL (Installation Control Library) for the latest consideration for HFS data sets.

**DIRECT
ACCESS
(DA-BDAM)**

DA (BDAM) data sets can be read or written directly, i.e., any record in the data set can be read or updated at any time. DA data sets are typically used by third-party software for database storage (such as SAS and some database vendors).

DUMP: All allocated tracks are dumped.

RESTORE: A physical restore will restore all allocated tracks in a DA data set. A logical restore will never reblock DA data sets, so it is used only for restores to unlike disk types.

Logical restore will not restore DA data sets which were created with absolute track addressing (OPTCD=A). Logical restore will rearrange the data blocks to fit on the new device type only for fixed-length (RECFM=F/FB) data sets; variable-length (RECFM=V/VB) or undefined length (RECFM=U) DA data sets will be restored by a physical ("track-image") restore and that restore will fail if the input data set has more data on a track than will fit on a track of the output disk.

A "track-image" restored DA data set may or may not be usable depending on the requirements of the program which accessed it; FDR has no way of knowing what those requirements are.

IAM data sets may be DA; see notes below. Other software vendors may use DA files for databases or other uses. Contact Innovation or the vendor for information on the use of FDR with these files.

ALLOCATION: if the output data set does not exist on the target output volume, FDR will allocate it. It will be allocated with the same number of tracks as the original unless the CYL=, TRK=, RLSE=, or %FREE= operands are specified (for a logical restore to an unlike device the space may be recalculated). The allocated space may be in fewer or more extents than the original, but the total must equal or exceed the number of tracks to be restored.

MULTI-VOLUME: Since FDR is a volume-oriented product, the pieces of multi-volume DA data sets will be backed up to separate backup files. During the restore, FDR will allocate each piece of the multi-volume data set on a separate volume, **so it must be restored to as many volumes as it was dumped from**. The data set will not be usable until all original pieces have been restored.

80.11 CONTINUED

UNMOVABLE PS, PO and DA data sets can be marked as unmovable in the Format 1 DSCB (via DSORG=PSU, POU or DAU). An unmovable data set can be restored ONLY to the same track addresses on the output volume as it occupied on the input volume. If the data set does not already exist on the output volume, and it was allocated as a single extent on the input volume, FDR will attempt to allocate space for it at the same locations. If it occupied more than one extent on the input volume or if the required space is not free on the output volume, it cannot be allocated. If the user pre-allocates space for an unmovable data set, it must be allocated at the correct locations; if any extent is not at the same location, message FDR111 message is issued, indicating where the allocations should be made. Unmovable data sets cannot be restored to an SMS-managed volume.

IAM DATA SETS IAM (Innovation Access Method) is a separate software product from Innovation which is a high-performance replacement for ICF VSAM KSDS (indexed) and ESDS (sequential) data sets. IAM data sets may be DSORG=DA (direct access) or DSORG=PS (physical sequential) with RECFM=F. In either format, FDR can properly restore them to a like or unlike disk, with no special considerations.

SAS DATA SETS SAS data sets created by SAS Version 5 or below are a common type of BDAM RECFM=U data set (SAS is a program product of SAS Institute). SAS data sets can be restored to the same disk type (physical restore) with no special considerations. A logical restore to a new disk type uses the track-image mode described above, so they can be restored only to disks with a larger track size; after a logical restore they can be used by SAS, but it may be necessary to specify "DEVFMT=olddevicetype" (e.g., DEVFMT=3380) on a SAS LIBNAME statement; a SAS PROC COPY should be done on them to reformat them properly for the new disk type. Because of this, SAS PROC COPY rather than FDRCOPY is recommended to move SAS Version 5 data sets to a new device type.

Sometimes SAS data sets are incorrectly marked as unmovable (DSORG=DAU), although SAS Institute assures us that there is no reason to do so and they are in fact movable. In that case the rules for Unmovable data sets will apply (see below), unless you correct the DSCB by turning off the Unmovable indicator. Contact INNOVATION if you need assistance restoring a SAS data set that was marked DSORG=DAU when dumped.

In SAS Version 6 and above, SAS files are now PS (sequential) with RECFM=FS (Fixed standard); FDR handles these files without special considerations.

MODEL DSCBS A model DSCB is a data set with no tracks allocated (zero extents). They are usually used to contain model DCB information for a GDG (generation data group). ABR uses a model DSCB on each volume to contain options and information about ABR's use of the volume.

DUMP: Since there are no tracks associated with it, only the DSCB for the model is dumped in the control records on the backup data set.

RESTORE: If necessary, the model will be allocated with no extents. The restore will update the model with the characteristics of the original. Note that GDG models are usually uncataloged data sets which must exist on the same volume as the catalog in which the GDG is cataloged, so it may be a mistake to restore or copy them to a new volume. If model DSCBs are moved with FDRCOPY MOVE, FDRCOPY will be unable to catalog them and so will not scratch the original model DSCB (unless NOCAT is specified).

The ABR model DSCB (FDRABR.Vvolser) will never be restored or moved since its information relates to the volume it exist on. Any attempt to select it results in a warning message. The ABR model can be created or changed by the ABR ISPF dialogs (Option A.I.8) or by program FDRABRM.

80.11 CONTINUED

**GENERATION
DATA GROUPS
(GDG)**

DUMP: Since individual GDG generations are simple non-VSAM data sets, they are dumped as described earlier based on their DSORG (usually PS).

RESTORE: GDG generations are restored according to the rules for their DSORG,

ALLOCATION: Each GDG generation has an absolute generation number (e.g., G0023V00) that increases with each new generation created. In normal use, they are usually referred to by relative generation number, relative to the most recently created generation. For example,

gdgname(+1)	creates a new generation
gdgname(0)	reads the most recently created generation
gdgname(-3)	reads the third oldest generation.

By default, restore of a generation will allocate and catalog it under its original absolute number. If the current generations in that GDG all have higher absolute numbers, and the GDG index is full, catalog management will uncatalog (and may delete) the absolute generation that is currently the lowest one in the GDG (even if it has a higher absolute generation number than the one being restored), and catalog the new one in its place. **This may cause unexpected data loss.** If multiple generations are being restored at one time, it may be that only the last one cataloged will be restored correctly. If it is your intention to restore the GDG to its state at an earlier time, you should DELETE all of the currently cataloged generations and let the restored generations get cataloged in their place.

GDGs can be renamed to a new relative generation number using the NEWNAME= or NEWINDEX= operands on the SELECT statement, e.g., NEWI=(+1) will restore a GDG as the next available generation number.

See [Section 70](#) for details on SMS-managed GDGs.

**STANDARD
USER LABEL
DATA SET
(SUL)**

Non-VSAM data sets may be allocated with standard user-labels (LABEL=SUL) which allocates an extra track for the storage of user labels containing application-dependent information separate from the data. FDR will allocate such data sets correctly; if pre-allocated by the user, LABEL=SUL must be specified. The label track will always be restored with a physical (track-image) restore even if the data is being restored logically.

IBM DB2

DB2 files are ICF VSAM clusters with a special format; they are non-indexed files with a 4K CISIZE and none of the usual internal VSAM indicators; DB2 manages all data in the file. They are allocated as LINEAR clusters. FDR recognizes them by the special format of their cluster and component names:

a.DSNDBx.b.c.I0001.Annn

where a, b, and c are user-chosen, x is either C for the cluster or D for the component, and nnn is a numeric sequence number. Remember that FDR always processes ICF clusters by the cluster name (the DSNDBC name for DB2), never by the component name.

FDR can dump and restore DB2 files just like any other ICF VSAM cluster, including moving or restoring to unlike disks, with these considerations:

- 1) if you are backing up active DB2 files, you should quiesce them or place them in read/only mode during the backup, using appropriate DB2 commands.
- 2) if a DB2 file was allocated by DB2 onto a volume in a DB2 storage group (different from an SMS storage group), it can only be moved or restored to another volume in that same DB2 storage group.
- 3) Since DB2 files must be recorded by DB2, and FDR has no way to update this information, DB2 files cannot be copied or restored to a new name; to do so makes them unusable by DB2. You must use DB2 utilities to copy DB2 files to new names and get them properly recorded.
- 4) If you restore a back-level copy of a DB2 file, you may need to run the DB2 REPAIR utility to correct discrepancies between the data and DB2 records.

If you restore a DB2 file from an FDR backup, you may be able to bring the database up to a more current level by using the DB2 RECOVER utility to re-apply updates from the DB2 logs.

DB2 manuals contain more information on using non-DB2 utilities such as FDR with DB2 data.

80.11 CONTINUED

ICF VSAM [Section 80.13](#) contains details on support for ICF VSAM clusters. In general you can dump clusters and restore them to like or unlike disks, or copy/move them to like or unlike disks. You cannot restore or copy/move multi-volume VSAM clusters to unlike disks.

ICF VSAM clusters can also be EF (Extended Format), as described below.

EXTENDED FORMAT (EF) Many SMS data sets can be Extended Format (EF) data sets. EF data sets currently include:

- striped sequential data sets. A striped data set is usually multi-volume, where record 1 is on volume 1, record 2 is on volume 2, etc., so that multiple records can be read or written concurrently.
- Compressed sequential data sets.
- Compressed VSAM KSDS clusters.
- VSAM clusters over 4GB in size.

FDR can allocate and restore EF data sets. EF data sets must be allocated to a SMS volume which is at a hardware level which supports Extended Format. It is your responsibility to insure that your SMS ACS routines will allocate the data set to an appropriate volume.

EXTENDED ATTRIBUTE (EA) Many SMS data sets can also be Extended Attribute (EA) data sets. EA data sets are used with the Network File System (NFS). When an EA data set is restored to a SMS volume, its extended attributes are restored.

VTOC DUMP: The VTOC will always be dumped as a data set by full-volume dumps and by ABR Volume Backups. All tracks are dumped. In addition, any full-volume or data set dump will dump all DSCBs relating to the data sets dumped into control records at the beginning of the backup data set.

RESTORE: Any data set restore requires that a DSCB exist in the output volume's VTOC before it can be restored; this can be created by FDR or pre-allocated by the user. At the end of the restore that DSCB will be updated with information from the saved DSCB from the backup data set. (See [Section 80.10](#) for details.) The VTOC cannot be restored as a data set; absolute track restore can be used to do restore the VTOC, but it is usually a mistake to try to do so. FDRCOPY will never copy or move the VTOC.

Note: if you need to expand the size of a VTOC, this can be done with COMPAKTOR using a COMPAKT-from-backup. [See Section 40.](#)

VTOC INDEX DUMP: If the volume has an indexed VTOC (SYS1.VTOCIX.xxxxxx), it will be dumped as a data set if selected, by full-volume dumps, and by ABR Volume Backups.

RESTORE: The indexed VTOC on an output volume is updated by allocation of data sets and can be reconstructed from the information in the VTOC by the IBM utility ICKDSF, so it is always a mistake to restore or copy it. Any attempt to select it results in a warning message.

80.11 CONTINUED

VVDS DUMP: The VVDS (VSAM Volume Data Set, SYS1.VVDS.Vxxxxxx) will be dumped as a data set if selected, by full-volume dumps, and by ABR Volume Backups. All tracks are dumped. Also, the VVDS records (VVRs for ICF VSAM and NVRs for non-VSAM SMS-managed data sets) are also recorded in control records at the beginning of the backup data set.

RESTORE: Allocation of other output data sets by FDR or by the user will create the VVR and NVR in the VVDS of the output volume. The VVDS will be automatically created by allocation if it does not exist on the volume). FDR will update the VVR/NVR information at the end of the restore. It is almost always an error to try to restore the VVDS as a data set; any attempt to select it results in a warning message. More details are in [Section 80.13](#).

ICF VSAM CATALOG [Section 80.13](#) contains details on support for ICF VSAM catalogs. FDRCOPY cannot be used to move or copy an ICF catalog, and an ICF catalog cannot be restored to a new device type. IDCAMS Export/Import or Repro may be used if it is necessary to move ICF VSAM catalogs to a new device type (unlike disk).

80.12 DATA MOVEMENT BETWEEN DIFFERENT DASD DEVICES

The FDR system can be used to move or restore data from one disk device type to another. This section describes some of the techniques and rules for doing so.

A disk device is characterized by the maximum number of bytes that can be written on a track, and by the number of tracks in a cylinder; these are called its "device geometry". FDR considers two disks to be "unlike" devices if either of those two values is different. For example, a 3380 and a 3390 disk have the same tracks/cylinder (15) but the bytes/track is different (47476 vs. 56664) so FDR treats them as unlike devices.

However, within a type of disk, such as 3380 or 3390, there are different models that have a different total number of cylinders. For example, 3390s have 1113 (3390-1), 2226 (3390-2), 3339 (3390-3) or 10017 (3390-9) cylinders. Some disk subsystems allow you to define disks of any arbitrary number of cylinders. Despite the difference in size, these are considered "like" devices for restore. The number of cylinders is also called the "density"; two disks with different density have the same geometry but different cylinder count.

**FULL-VOLUME
RESTORE**

An FDR or ABR full-volume restore (RESTORE TYPE=FDR) can be used only to restore to like devices, disks of the same geometry which may differ in size (total number of cylinders). Since full-volume restore rewrites the exact images of the original tracks, it cannot restore to an "unlike" device since the data would not be formatted properly for the different device geometry.

Restores to a volume with a larger number of cylinders than the original disk are fully supported by FDR and ABR. Since all of the tracks being restored exist on the output disk, they can all be restored properly. All that is required is to make the excess tracks available as free space for allocation of new data sets.

Restores to a volume with a smaller number of cylinders can also be done, as long as the original volume has no data sets allocated on cylinders which do not exist on the output disk. To enable this restore, specify the PROT=NONE operand on the RESTORE or COPY TYPE=FDR statement. Tracks that no longer exist will be removed from the free space on the volume. However, COMPAKTOR is a preferable method for doing this conversion, as described later in this section.

To update the free space, at the end of the restore, FDR will set the DOS flag in the Format 4 DSCB in the VTOC, telling MVS that the Format 5 (free space) DSCBs are invalid and must be rebuilt. Then a dummy data set will be allocated causing this to occur. The dummy data set's name is

FDRABR.VvvvvvZ

where "vvvvvv" is the disk volser. However, this allocation is designed to fail and will cause a IEC614I RC 16 message in the joblog; this is expected and is sufficient to cause the Format 5 DSCBs to be rebuilt to reflect the new size of the volume. Your security system, if any, must allow the allocation of this data set for the free space update to occur. This allocation will not work on SMS-managed volumes; you will receive message FDR341 COMP=192 and must manually allocate a data set on the volume to update the free space.

If the original volume had an active indexed VTOC, it will be disabled after the restore and must be rebuilt by the user (see example in [Section 10.11](#)). This is especially important on SMS-managed volumes.

After a full-volume restore, all data sets, including the VTOC, indexed VTOC, and VVDS will have the same size and position as they had on the original disk.

**FULL-VOLUME
COPY**

An FDR full-volume disk-to-disk copy (COPY TYPE=FDR) can also copy between different densities of the same disk type. The rules are the same as for a full-volume restore.

80.12 CONTINUED

COMPAKTOR A COMPAKTion from a backup (not a "Fast COMPAKTion") is essentially a full-volume, track-image restore, so COMPAKTOR is also limited to restoring to like disks with the same geometry.

A COMPAKTOR restore to a different density disk has several advantages over an FDR full-volume restore:

- Since COMPAKTOR can move data sets during the restore, it will automatically move data sets to fit in the new space.
- When restoring to a smaller disk, it will relocate all data sets to fit unless the total tracks required exceeds the capacity of the output disk, or there are unmovable data sets which are located beyond the end of the new disk.
- If there is insufficient space, the SELECT statement operand SCRATCH=YES can be used to eliminate selected non-VSAM data sets until the remaining data sets do fit (for non-SMS-managed volumes only).
- COMPAKTOR can move the VTOC to an appropriate location on the output disk. It can also expand the VTOC, which may be required on a larger disk since more data sets may be allocated there.
- COMPAKTOR can move (but not expand) the indexed VTOC and VVDS.
- COMPAKTOR will automatically rebuild the Format 5 or Format 7 DSCBs indicating the proper free space for the size of the output disk.
- If the volume had an active indexed VTOC, COMPAKTOR will automatically rebuild it.
- COMPAKTOR supports output disks with a non-standard number of cylinders, by honoring the cylinder count in the VTOC.

Since a Fast COMPAKTion (FASTCPK) always reorganizes one online disk, unlike device is not an issue.

DATA SET RESTORE A DSF or ABR data set restore can be done to a "like" device or, in most cases, to an "unlike" disk device.

When restoring to a "like" disk device, all densities and models of the disk are treated the same, since the disk geometry is the same. A "physical" track-image restore is done (unless BLKF= is specified to reblock PS data sets). If the output data set must be allocated, there must be sufficient free space on the volume to contain the data set, for example, a data set which occupied an entire 3390-3 cannot be restored to a smaller 3390 and a 100-cylinder data set must be restored to a volume with at least 100 free cylinders (in 5 or fewer extents), unless the space requirement is overridden by the user. Details on how physical device restore handles various data set types are in [Section 80.11](#).

Data set restore has no facility for "breaking apart" a single data set onto multiple volumes. Likewise, a data set which was in multiple parts on several original disks cannot be combined into one piece on a larger disk; it must be restored to as many disks as it was dumped from. Data sets which were marked as "unmovable" and were allocated in the higher tracks of a larger disk cannot be allocated on a smaller disk, since those tracks are not available.

When restoring a data set to an "unlike" disk, since their device geometry differs, a "logical" restore will be done. Although the original backup is still in track-image format (FDR has no need to do a special logical backup), the data will be extracted and reformatted for the new device type. The amount of space required for each data set will be adjusted so that the output data set occupies about as many bytes as the data set did on its original disk. Details on logical restore of various data set types are also in [Section 80.11](#).

A logical restore may also be done if data set reblocking is requested, or for data set types requiring special processing.

When restoring to a new disk volume serial, DSF and ABR can recatalog the data set to the new disk, if the RECAT (for non-VSAM data sets) and VRECAT (for ICF VSAM clusters) operands were specified. VRECAT will delete the original cluster if it still exists.

80.12 CONTINUED

**DATA SET
COPY/MOVE**

FDRCOPY can copy or move data sets to "like" and "unlike" devices, under the same rules as for data set restore (FDRCOPY is essentially a track-image dump and a physical or logical restore in the same step). Although FDRCOPY is a very fast way of moving data sets to a new disk, and can run without operator intervention (tape mounts), no backup is created by FDRCOPY.

**CONVERTING
TO LARGER
DISKS**

Often, an installation will replace disks of one capacity with "like" disks of a larger capacity (such as replacing 3390-2 disks with real or RAID-emulated 3390-3 disks). The lower density disks may need be combined onto a smaller number of high-density disks.

If possible, one of the original disks can be directly converted to a larger disk under its original volume serial. This can be done with an FDR or ABR full-volume restore, an FDR full-volume COPY, or COMPAKTOR. Since data sets restored by a full-volume operation are NOT recataloged, this restore MUST be done using the original volume serial (CPYVOLID=YES). COMPAKTOR is preferable for this operation since it can expand and relocate the VTOC. Examples of full-volume conversion can be found in [Sections 10.11, 10.12, and 40.23](#). Data sets from other original volumes can be added to the larger disk with DSF or FDRCOPY as described below.

If the full-volume operation is not possible (if, for example, the new disks must have new volume serials), then ALL of the original data sets must be moved with DSF or FDRCOPY.

To move data sets to the larger disk:

- FDRCOPY can easily move all data sets on a volume to a new volume. This is the recommended method. An FDRCOPY MOVE will automatically recatalog data sets to the new volume. Examples can be found in [Section 21.11](#). See [Section 80.13](#) for considerations on moving multi-volume VSAM clusters.
- DSF can restore all data sets from an FDR full-volume backup or DSF backup of the original disk volumes. RECAT and VRECAT should be specified so that all data sets are cataloged on the new volume. [Section 20.12](#) contains an example.

**CONVERTING
TO SMALLER
DISKS**

It may be necessary to fall-back or convert from high-density disks to lower-density disks of the same type (such as from triple density 3390-3 volumes to double density 3390-2 volumes).

If the data sets on the larger disk will all fit on the smaller disk, then a full-volume restore or COMPAKTion may be used to do the conversion, as described earlier in this section.

In most cases the data will not fit on a single output disk, and must be spread over multiple output disks. The DSF and FDRCOPY techniques just described for converting to hi-density disks can also be used for conversion to lower-density disks. The NVOL= operand on SELECT statements for DSF, ABR, or FDRCOPY must specify multiple disk volumes, e.g.,

```
NVOL=(PROD01,PROD02,PROD03)
```

or

```
NVOL=PROD*
```

which will spread the selected data sets over the indicated volumes. Examples can be found in [Sections 20.12 and 21.11](#).

**CONVERTING
TO NEW DISK
TYPES**

Since full-volume operations do not work with "unlike" disk devices, only data set operations can be used to convert data sets to a new disk device type (such as 3380 to 3390). The DSF and FDRCOPY techniques just described for conversion between models of "like" disks can be used for "unlike" conversion.

However, there are special considerations and restrictions on certain data set types; the details are in [Section 80.11](#).

80.13 ICF VSAM SPECIAL CONSIDERATIONS

This section explains some of the special considerations that a user should be aware of when dumping or restoring VSAM clusters which reside in ICF catalogs (almost all VSAM clusters will be in an ICF catalog).

There are special rules and considerations for multi-volume ICF clusters. If multi-volume clusters are involved, be sure to read *Multi-Volume VSAM Considerations* at the end of this section.

WHAT FDR DUMPS FDR, DSF and ABR will backup ICF VSAM Clusters. All of a cluster's associated components which reside on a single volume will be dumped. This includes the index, data, alternate indexes, key ranges, etc. In addition, FDR will save the VVR information contained in the VVDS for each component in control records at the beginning of the backup. Information on PATH names of alternate indexes and aliases of user catalogs is also extracted from the catalog and saved on the backup. The only information FDR does not preserve is other information exclusively maintained within the catalog, including passwords, RACF indicators and Candidate Volumes. These fields are not recorded by FDR and will not be restored. On non-SMS-managed volumes, the expiration date is also only in the catalog and is not preserved; on SMS volumes, the expiration date of ICF clusters is preserved.

FDR dumps VSAM files by track image, which is the same technique used for every other type of file. FDR is not access method oriented. If logical errors exist within the file, they will not be detected by FDR. However, certain types of errors within the VVDS data set will be detected and reported.

For DSF and ABR data set backups, SELECT statements should select ICF VSAM clusters only by the base cluster name, never by individual component names or alternate index cluster names. Even if a volume contains only an alternate index (AIX), the base cluster name must be used, not the AIX cluster name. However, COMPAKTOR SELECT statements, which may be used to position ICF VSAM components, will accept ONLY component names.

VSAM LAST REFERENCE DATE VSAM OPEN will update the DSCB with the update flag (for OPEN OUTPUT or UPDATE) and the last reference date (for any OPEN), but only in the DSCB of the data component of the base cluster and only on the first volume of a multi-volume cluster. They are not set in the DSCBs of index components or alternate index components. They are not set at all in the base data component for a keyrange KSDS.

ABR VOLUME BACKUPS ABR incremental Volume Backups will select ICF VSAM clusters. If the update flag is on for any component of a cluster, ABR will backup all of the associated components which reside on the volume being dumped. Since the update flag is never set on secondary volumes, when ABR finds that a volume contains only secondary pieces of a cluster it will always back them up. ICF VSAM can be totally excluded only by a statement such as:

EXCLUDE ALLDSN,DSORG=EF

ABR APPLICATION BACKUPS ABR Application Backups will backup ICF VSAM clusters. Since the catalog is normally used to locate data sets for Application Backup, all components of the selected clusters will be backed up from cataloged volumes. The Application Control File will record all cluster and component names.

80.13 CONTINUED

ABR ARCHIVE AND SUPER-SCRATCH ABR can Archive or Superscratch ICF VSAM Clusters. VSAM clusters can be selected manually (SELECT statement) or by automatic criteria such as ADAYS=. For the options ADAYS= or ADATE=, ABR will only use the last reference date stored in the data component of the base cluster. However, if ABR finds that any part of a cluster being examined has been archived from another volume, it will archive the pieces on the current volume.

The operands IFNOTCAT and MAXGDG= do not apply for VSAM clusters and are ignored. The EXPIRED operand only works on SMS-managed clusters. If the SIZE= operand is specified, this size is checked against each of the components; if any component exceeds the size specified, the entire cluster will be archived.

The operator will not be given message FDRW23 on unexpired VSAM Clusters. If the VSAM Cluster is not expired, ABR will not scratch the cluster unless VEXPD=NONE or EXPD=NONE is specified on the DUMP control statement.

The ARCHIVE Control file will record the cluster name and each of its component names. ABR will print an FDR314 message for each of the components selected, and when the last component for the cluster is printed, it will print an FDR314 message for the cluster name. However, the cluster is not scratched until it has been backed up from each volume that it exists on; all of those volumes must be processed in the same ABR Archive step.

VSAM RESTORE RULES For FDR and ABR full-volume restores, the VVDS and all VSAM components are, of course, restored. It is the user's responsibility to be sure that the clusters are properly cataloged; no cataloging will be done by a full-volume restore.

DSF and ABR data set restores can restore ICF VSAM files from data set or full-volume backups. The base cluster name, not individual component names, is always used to select ICF VSAM. Once selected, all of the components of a cluster which exist on the backup (from one single disk volume) will be restored; this includes the data, index, alternate indexes and key range components. An alternate index cluster cannot be restored individually unless it resides on a volume by itself (see *Multi-Volume VSAM Considerations*), and even then it is selected using the base cluster name, not the AIX cluster name.

Data set restores can be "physical" or "logical". A physical restore is a "track-image" restore and is used when the cluster is being restored to the same disk device type as the original cluster resided on; the tracks of the data set look exactly as they did when dumped, and the cluster will not be reorganized or rearranged in any way. A logical restore is used when restoring a VSAM cluster to a different device type (a "unlike" restore); the control intervals in the component will be rearranged to fit on the new device type, but the cluster will still not be reorganized. Use FDRREORG to reorganize VSAM KSDS clusters.

Logical restore may change cylinder allocated clusters to track allocated; this is normal and causes no performance problems. Logical restore is sometimes used for KSDS clusters with the IMBED option, even when restoring to a like device.

NOTE: Logical restore for VSAM supports all types of ICF VSAM clusters. For current support and considerations for ICF VSAM to unlike devices, see member VSAMUNLK in the Installation Control Library (ICL).

WARNING: Because of the way that logical restore processes the control areas of ICF VSAM clusters, clusters can generally NOT be restored, copied, or moved to a disk with a smaller track size (such as 3390 to 3380) unless it's CA size (tracks per Control Area) is less than a cylinder. However, a cluster which was moved to a disk with a larger track size by FDR or ABR can be moved back to its original disk type if it has not been reDEFINED since the move.

80.13 CONTINUED

**VSAM
RESTORE
RULES
(continued)**

After the individual components have been restored, the VVR information in the VVDS data set will be updated for each component. Information exclusively maintained in the VSAM catalog will not be updated. This includes passwords, RACF indicators and candidate volumes; it also includes the expiration date for ICF clusters on non-SMS-managed volumes. This is not a problem if the data set is being restored back into the original VSAM cluster. If you pre-DEFINE the cluster for the restore, you should specify this information in the DEFINE. If FDR DEFINES the cluster, you should re-establish this catalog information with the IDCAMS ALTER command after the restore. If the cluster was protected by a RACF discrete profile, that protection must be re-established in the catalog with a RACF command.

When an alternate index (AIX) is allocated by FDR, all PATH names associated with that alternate index are also defined. However, PATHs which relate only to the base cluster (similar to aliases for non-VSAM data sets) are not defined; you must manually do a DEFINE PATH for them.

The restore printout will report on each component name restored, giving its associated cluster name.

**FDRCOPY of
ICF VSAM**

An FDRCOPY disk-to-disk data set copy or move ([see Section 21](#)) is essentially a simultaneous dump and restore, so the preceding rules for backup and restore of ICF VSAM apply. However, since the input cluster still exists during the copy, information only in the catalog (passwords, expiration date, etc.) will be copied to the equivalent output cluster. If the input cluster was protected by a discrete RACF profile, a discrete profile will be created for the output cluster using the input profile as a model.

Since two ICF clusters by the same name cannot exist in the same catalog, an FDRCOPY COPY of an ICF VSAM cluster must be to a new name unless the new cluster is directed to a new catalog; these options are described in following paragraphs. A MOVE of an ICF VSAM cluster is accomplished by assigning temporary names to the output cluster and its components; they are renamed after the input cluster is deleted. [See Section 21](#) for details.

FDRCOPY supports both "physical" and "logical" copies (as detailed under *VSAM Restore Rules*), so clusters can be copied or moved to "unlike" disk devices.

**NEWNAME
RESTORE**

An ICF VSAM cluster can be restored or copied/moved to a new name (NEWNAME=) or new group name (NEWGROUP=) or with new or replacement index levels (NEWINDEX=). ICF catalogs and SYS1.VVDS data sets cannot be renamed so they cannot be copied or moved, only restored..

If NEWNAME= is specified, this name will be used as the new cluster name. If the named output cluster already exists, a LOCATE will be done to determine the associated component names for that cluster, matching each type of component (data, index, alternate index) with the corresponding component from the backup. If NEWNAME= is specified for clusters which contain more than one alternate index, only the base cluster will be correctly restored.

If the cluster named by NEWNAME= does not exist it will be allocated by FDR, allowing VSAM to generate default names for the components. FDR will use the generated names as implied "newnames" for the components. A VSAM cluster which contains alternate indexes cannot be allocated if NEWNAME= is specified.

INNOVATION recommends the use of NEWGROUP= or NEWINDEX= instead of NEWNAME= if the output cluster must be allocated.

If NEWGROUP= or NEWINDEX= is specified, FDR will apply the group name or indexes to the cluster name and all of the associated components, including any alternate index clusters and their components. If any of the cluster or component names are too short or contain too few index levels, or if the generated new names are duplicates of one another, the restore will fail. If the new cluster name does not exist, the cluster will be allocated using all of the new component names. If the new cluster is pre-allocated, then all of the existing component names must match the new names generated by FDR.

80.13 CONTINUED

**ALLOCATING
VSAM
CLUSTERS**

If the components of the output cluster are found on the output disk, such as when restoring the backup copy of an existing cluster, FDR will simply overlay the contents of the components; no allocation is required. However, a logical (unlike) restore of ICF VSAM cannot restore to a pre-allocated cluster.

If the output cluster name is not found on the target disk volume, the cluster will be allocated (see *Multi-Volume VSAM Considerations*). The cluster will be allocated so that each component has a primary allocation equal to the total amount of space the data set occupied when it was dumped, so a component that was in multiple extents when dumped may occupy fewer extents when restored. The primary allocation of the data component of the base cluster may be overridden by the CYL= or TRK= operands of the SELECT statement, but the value specified must be equal to or greater than its original allocation. On track or record allocated clusters, VSAM will round up to a multiple of the number of tracks in the CA. When doing a "logical" restore or copy/move to a different disk type, the allocation quantity will be adjusted.

The ICF catalog into which the new cluster will be defined depends on the ICFCAT= parameter on the RESTORE statement:

- ICFCAT=ORIGINAL: the cluster's original catalog name is used. If that catalog does not exist on the system where the restore is run, the allocation will fail with message FDR157 COMP=0004 CODE=00120; use ICFCAT=ALIAS or STEPCAT to direct the cluster to a different catalog. ICFCAT=ORIGINAL is ignored when restoring a cluster to a new name; ALIAS will be forced instead.
- ICFCAT=ALIAS: FDR will attempt to locate the user catalog associated with the new cluster name in the master catalog; if found, this catalog will be used; if the alias does not exist in the master, the STEPCAT (if present) or the master catalog will be used.
- ICFCAT=STEPCAT: may be specified to force the cluster to be defined in the STEPCAT (if present) or master catalog.

ICFCAT=ORIGINAL is the default when restoring a cluster to its original name. When restoring to a new name, ICFCAT=ALIAS is the default.

Since ICF VSAM clusters MUST be cataloged at the time they are allocated, the define of a cluster will normally fail if the cluster name or any component name already exists in the catalog in which FDR is attempting to catalog the cluster. This might be the case if the cluster currently exists on another volume, or if you are restoring a cluster which does not exist on its target volume but is still cataloged. The VRECAT operand on RESTORE commands allows FDR to allocate a cluster even if it is in the catalog. VRECAT will attempt to DELETE the cluster; if that fails, a DELETE NOSCRATCH is attempted.

WARNING: To avoid leaving uncataloged VSAM components on other volumes, VRECAT will SCRATCH the cluster if it exists on another volume.

80.13 CONTINUED

**USER
ALLOCATED
VSAM
CLUSTERS**

You may manually pre-allocate output ICF VSAM clusters using IDCAMS DEFINE prior to a restore or copy/move. However, user-allocated clusters cannot be used for a logical (unlike) restore, and usually cannot be used for multi-volume clusters or for KSDS clusters with keyranges or the IMBED option if they have gone into additional extents.

The DEFINE must result in a cluster with the same options and characteristics as the VSAM cluster to be restored. This includes CISIZE, imbedded index (IMBED), physical blocksize, type of VSAM file (ESDS, KSDS, etc.) and so forth. FDR may not be able to restore the cluster if some of the characteristics are different. Message FDR152 will be issued if any required characteristics do not match, detailing the comparison which failed.

The PRINT TVTOC report ([Section 53.09](#)) can provide an IDCAMS-style listing of the characteristics the cluster had when it was backed up.

Cylinder-allocated VSAM components must be re-allocated in cylinders, while components originally allocated in records or tracks must be re-allocated in tracks. The CYLINDERS or TRACKS keyword for each component must specify the total number of cylinders or tracks that the component occupied when it was dumped. Even if the space parameters were originally specified on the CLUSTER level, they should now be specified on each component to get the space allocated correctly.

For components allocated in tracks, a secondary allocation quantity equal to the CA size (tracks-per-Control Area) of the component being restored should be specified to set the CA size correctly, even if the original secondary allocation quantity was zero. For a KSDS with the IMBED option (imbedded index), subtract 1 from both the primary and secondary track quantity for the data component since VSAM will add it back.

FDR will reset the primary and secondary allocation to their original values after the restore.

The new allocation can be larger than the original. However, the result must be a multiple of the CA size of the original cluster. An error message will result if the allocated space is too small for any component, showing the required number of tracks.

Example: The VSAM data component was allocated with a primary of 3 tracks and a secondary of 3 tracks. The data set acquired secondary allocation and now occupies 9 tracks. The number of tracks per CA is 3.

TRACKS(9 3)

should be specified for the data component but if it has the IMBED option use

TRACKS(8 2)

If a VSAM cluster being restored has the original cluster name but the component names are different from the original cluster, specify the same cluster name in both the DSN and NEWNAME operands. This forces FDR to locate the new component names from the catalog.

You are trying to restore a KSDS with the IMBED option (imbedded index) or keyranges and a component has expanded into additional extents, it may not be possible for you to pre-allocate the cluster since the arrangement of the component RBAs will not meet FDR's requirements. However, FDR can allocate the cluster properly, so FDR will issue messages telling you to scratch the cluster and try the restore again.

Instead of trying to DEFINE the cluster yourself, it is usually easier and faster to let DSF or ABR do the DEFINE. For logical (unlike device) restore or copy/move, you must let DSF or ABR define the cluster.

80.13 CONTINUED

**ICF CATALOG
RESTORE**

DSF and ABR can dump and restore ICF catalogs to a like disk type. ICF catalogs can be restored without restoring or affecting the clusters cataloged within them. After a restore, an IDCAMS DIAGNOSE should be run against the catalog to detect any inconsistencies between the catalog and the VVDS data sets associated with it. This will detect clusters which are in the catalog but were scratched since the backup. These entries can be deleted using IDCAMS with a DELETE NOSCRATCH Command. A DIAGNOSE of all the VVDSs which reference this catalog will detect clusters which were added since the backup. An IDCAMS DEFINE RECATALOG command will re-establish these clusters in the catalog; information exclusively maintained in the catalog will not be re-established automatically, but must be specified on the DEFINE RECATALOG command.

This is a significant advantage over other methods of restoring these catalogs which require the individual clusters to be restored in order for the catalog to be recreated.

ICF catalogs can be restored to a new volume. FDR will alter the self-defining records within the catalog to reflect the new volume serial. However, ICF catalogs can only be restored to the same disk device type they were dumped from.

ICF catalogs cannot be restored to a NEWNAME. The reason for this restriction is that the clusters cataloged within this catalog, including its own self-defining record, would not know the NEWNAME. If a catalog is re-allocated, VSAM may assign a new name to the index component of the catalog. FDR is designed to recognize this change without the use of NEWNAME.

DSF and ABR will allocate an output ICF catalog if it does not already exist on the target volume, automatically connecting it into the master catalog of the system it is restored on and the VVDS on the volume it resides on. You cannot restore an ICF catalog to a new volume while the original catalog is known to the master catalog, because this would create a duplicate entry in the master catalog. You must use IDCAMS to either DELETE the original catalog (possibly specifying RECOVERY), or do an EXPORT DISCONNECT.

Because of the above considerations, FDRCOPY cannot copy or move an ICF catalog.

When an FDR restore DEFINES an ICF catalog, it will also define any aliases that the catalog had at the time it was dumped.

You should not restore a catalog in the same jobstep that restores any data sets that are cataloged in that catalog. For example, if the restore does a DEFINE for a data set, the new catalog entry will be wiped out when the catalog is restored.

**VVDS DUMP
AND RESTORE**

DSF or ABR can backup the SYS1.VVDS.Vvvvvvv data set. If a restore of this data set is to be done, it should be done with care. This data set is, in effect, the VTOC for the ICF VSAM files on this volume (and on SMS-managed volumes, for non-VSAM data sets as well). If a VVDS is restored, there exists the potential that the information contained in the VVDS, the VTOC and the catalog will be out of synch. An IDCAMS DIAGNOSE should be run after a VVDS restore, on the VVDS and all the catalogs which use the volume, which will detail most of the inconsistencies between the VVDS, the VTOC and the ICF catalogs. Any errors must be manually corrected.

To ensure that the VVDS is not accidentally restored, DSF and ABR will not restore the VVDS unless the fully qualified name is specified (DSN=) and PROT=NONE is coded on the RESTORE statement. The VVDS should be the only VSAM cluster restored in this execution. If the VVDS is included as part of an ALLDSN or DSN=mask restore, it will be bypassed with a warning message.

80.13 CONTINUED

**VVDS DUMP
AND RESTORE
(continued)**

The SYS1.VVDS data set cannot be restored to a newname, and it must be restored to the same tracks that it was dumped from. The reason for these restrictions is the presence of self-defining records within the VVDS and every catalog which references it. So if the VVDS has been deleted or the volume reinitialized, it is not possible for FDR to allocate and restore a usable VVDS (except by full-volume restore). If the VVDS has been lost, you will generally have to recover a backup copy of every VSAM cluster on the volume (non-VSAM data sets as well on a SMS-managed volume).

If the volume serial number is changed on a volume containing a VVDS, all of the VSAM clusters on that volume will be inaccessible. The FDR system will not be able to dump and restore them by cluster name.

FDRCOPY cannot copy or move a VVDS.

Since FDR updates the VVDS during the allocation and restore or copy/move of an ICF cluster, it is almost always a mistake to explicitly restore the VVDS, except during a full-volume restore. Contact INNOVATION for advice before attempting to restore the VVDS.

**MULTI-
VOLUME
VSAM
CONSIDER-
ATIONS**

Multi-volume VSAM Clusters are VSAM clusters having components which reside on multiple volumes. There are four types of multi-volume clusters:

1. A KSDS with the index component on one volume and the data component on another
2. A KSDS with an alternate index on a different volume
3. A key range KSDS with the keyranges on multiple volumes
4. Any type of cluster where any component, including a single keyrange component, has expanded to multiple volumes

In addition, there are two types of single-volume VSAM clusters which, because of their special allocation requirements, are treated like multi-volume clusters during allocation:

1. A "split imbedded index" KSDS, that is, a KSDS with the IMBED option, where the index component has expanded into additional extents.
2. Any keyrange KSDS.

Most of the comments in this section do not apply to these clusters, except for the explanation of allocation.

FDR and ABR full-volume operations can always handle multi-volume VSAM clusters, as long as all volumes involved are dumped and restored at the same time, and all volumes are restored to their original volume serial numbers. DSF, FDRCOPY, and ABR data set operations can dump, restore and allocate multi-volume clusters, but there are some special rules and considerations:

VSAM OPEN sets the update indicator and last reference date only for the data component on the first volume for the base cluster. Key range clusters never have indicators set even on the first volume. This presents a problem for ABR in that incremental backups cannot determine if the components residing on the remaining volumes need to be dumped. On incremental backups, ABR will always dump these multi-volume components except for an alternate index residing on a volume by itself.

Likewise, Archive and Superscratch cannot determine the last reference date when processing other than the first volume of the base data component. Archive may select a cluster from that first volume based on last reference date, but will not select the same cluster from any other volume based on date. Other Archive criteria, such as DSN= will select a cluster from all of its volumes. It is possible to select multi-volume clusters by last reference date by using FDREPORT as a pre-processor to ABR; if you need to Archive VSAM by ADAYS= and multi-volume VSAM might be involved, see the \$\$VSAM member in the FDR ICL (Installation Control Library).

80.13 CONTINUED

**MULTI-
VOLUME
VSAM
CONSIDER-
ATIONS
(continued)**

If Archive selects a multi-volume cluster from a volume, all components of the cluster on that volume will be backed up and recorded in the Archive Control File; however, Archive will not delete that cluster until it has been Archived from all its volumes. It is a requirement that the cluster must be Archived from all of its volumes in the same ABR Archive Backup step; otherwise ABR cannot tell when all pieces of the cluster have been processed (SELECT CATDSN= is a way to insure that all volumes are processed by ABR). RECALL=YES is supported for multi-volume clusters; when it is finally deleted by ABR, it will be recataloged as a multi-volume non-VSAM data set with the ABR RECALL indicators.

Superscratch (DUMP TYPE=SCR) will delete a multi-volume cluster from all of its volumes if any component meets the selection criteria on any processed volume, except that the last reference date will be checked only on the first volume of the base data component. As long as that base volume is among those processed, the cluster will be scratched.

Note on Archive and Superscratch: if a multi-volume cluster is not cataloged in the standard catalog search order (meaning that an alias in the master catalog points to the user catalog containing the cluster) and a STEPCAT is not provided to point to the right catalog, ABR cannot determine that the cluster is multi-volume. ABR will attempt to DELETE the cluster, but it will usually fail.

If you have backed up a multi-volume cluster and wish to restore it, and the cluster still exists on disk, you can simply restore back on top of it without deleting it, just like you can for single-volume clusters. However, the cluster must be in essentially the same state it was in when backed up; it must not have expanded to new extents, and must not have been deleted and redefined. FDR will simply overlay the extents of the cluster, and update the VVDS information appropriately. You must be sure that each piece of the cluster is restored back to the volume it was dumped from; this is automatic with ABR and under user control with DSF. Rules for restoring multi-volume clusters are explained later in this section.

If the cluster has been deleted from disk (or you are restoring or copying to a new name), it is usually not possible for you to pre-allocate multi-volume VSAM components in a way which is acceptable to FDR. It is much easier to let FDR allocate the multi-volume cluster.

For program FDRABRUT (Remote Queue) to support multi-volume VSAM clusters for the BACKUP or ARCHIVE commands, the DISKUPDATE=NO option must be enabled, causing the remote queue data set to be used.

80.13 CONTINUED

**ALLOCATING
MULTI-
VOLUME
VSAM**

In most cases of multi-volume clusters, including those with multi-volume components, FDR can allocate and catalog the cluster if it does not exist on the output disk, with these special considerations:

- a restore or copy/move is still a volume-by-volume operation. FDR will process one backup data set or one input disk at a time, and will attempt to allocate and restore only the components or parts of components on that backup or disk. However, FDR recognizes that the cluster is multi-volume and uses a unique technique for allocating those components. If the multi-volume cluster must be allocated by FDR, the volumes must be done one at a time; you can process all of the volumes in one job or jobstep, but you cannot run multiple restore jobs to restore in parallel. It does not matter what order the volumes of a multi-volume cluster are processed in.
- As it proceeds, the components on each output disk will be accurately allocated and restored as VSAM. However, until the final volume of the cluster is restored, the cluster will be cataloged as non-VSAM, with special indicators in the catalog entry. Once that last volume is processed the cluster will be recataloged properly as VSAM. So, the cluster will not be usable in any way until all pieces are successfully restored.
- If you display the catalog entry of a multi-volume cluster before FDR completes that last volume (e.g., with LISTCAT), it will appear as non-VSAM. The first volume serial will be #####Vx where "x" indicates the type of cluster; the other volsers will be the volumes where FDR has completed the restore. If you see this type of catalog entry, it is not abnormal, but simply means that there are more pieces of this cluster which must be restored or copied before it is usable. Another indication of multi-volume allocation status is the FDR311 message printed when FDR restores data sets; for a multi-volume VSAM cluster, that message will say ALLOCATED but not CATALOGED until the last volume of the cluster is processed (if multiple components exist on that last volume, only the last component says CATALOGED). Once CATALOGED appears, the cluster is usable.
- When allocating a cluster with multi-volume components, FDR must be able to allocate those components on the same number of volumes they originally occupied. If an output volume cannot be selected which does not already contain a part of the component, the operation will fail (but it can be rerun if another output volume can be provided). The rules for selection of output volumes depend on which program you are executing (FDRDSF, FDRCOPY, or FDRABR) and are detailed in [Section 20.07](#), [21.01](#) and [50.08/51.08](#) respectively. Providing a sufficient choice of volumes may be a matter of restoring to the original volumes, providing proper DD statements, providing an NVOL= operand, or setting up the ABR RESTORE ALLOCATION LIST to define volume pools. If the rules provide a choice of volumes, the allocation of a component or partial component may be tried on a number of volumes until one is found where it is successful.
- FDR cannot allocate an alternate index (AIX) which is itself on multiple volumes. However, the base cluster will be restored and the alternate index can be redefined and rebuilt with an IDCAMS BLDINDEX if necessary.
- Even a single-volume AIX cannot be allocated unless the base cluster has been previously restored or DEFINED, or unless the backup of the volume containing the AIX also contains the LAST piece of the base cluster to be restored. If you are doing a FDRDSF restore of a cluster with one or more AIXs on separate volumes, you must either arrange the restore so that the AIX volumes are restored after the volume containing the base cluster components, or re-restore the volumes containing the failed AIXs. ABR restores (including restores from volume backups, archive backups and application (FDRAPPL) backups), will automatically re-restore any AIXs which failed because their base cluster was not completely restored during the first restore attempt.
- If a component of a multi-volume cluster being allocated was originally in multiple extents, it will be allocated with the same number and size extents on the output volume.

80.13 CONTINUED

**RECOVERY
FROM
RESTORE
ERRORS**

Please read this carefully! FDR may successfully complete the restore of part of a multi-volume cluster from one or more input volumes, and then fail during the restore from another input volume; common causes include an insufficient number of output volumes and lack of free space on the output volumes.

When a restore error occurs on a multi-volume cluster, FDR will cleanup the pieces it tried to restore on the current output volume but will not cleanup the pieces of the cluster which were previously restored. They will be left on their volumes and the special non-VSAM catalog entry will be left.

If it is possible to correct the error, such as by choosing a new output volume with sufficient free space, you can restart the restore **but you must re-restore only from the input volume that cause the failure**. If the failure occurred on multiple input volumes, re-restore all of them.

However, if you want to re-restore the entire cluster from all volumes, you must manually cleanup the partially restored cluster before attempting the entire restore again. If you do not do this, various errors can occur.

To delete the non-VSAM left in the catalog, use the IDCAMS command:

```
DELETE clustername NONVSAM NOSCRATCH
```

Note that you may require special RACF authorization to issue a DELETE NOSCRATCH.

To delete the components or partial components which were successfully restored, use the IDCAMS command:

```
DELETE componentname VVR FILE(DD1) CAT(catalogname)
```

This must be executed against each component on each volume where the restore was successful. The FDR311 messages in the restore listing will show you which components were restored to which volumes. For a KSDS and VRRDS, part of the index component may have been restored to the same volume as part of the data component; you must issue two DELETE VVR commands to delete both components.

See the *IBM Access Method Services* manual for details and examples for these DELETE commands.

80.13 CONTINUED

**RESTORING
MULTI-
VOLUME
VSAM**

To restore a multi-volume VSAM cluster with FDRDSF, supply a TAPEX DD statement for the backup of each volume on which the data set resided, and code one SELECT statement specifying the cluster name. Output volumes can be specified by DISKx DD statements or by NVOL= parameters; if neither is specified DSF will restore back to the original volumes or, if the cluster currently exists on disk, to the volumes in the catalog.

Multi-volume clusters can be copied, but not moved, with FDRCOPY. Since FDRCOPY processes one input volume at a time, and the original cluster still exists, the special allocation techniques just described will not work unless a new name is given. So COPY requires a new name, which should be specified using the NEWGROUP= or NEWINDEX= operands (not NEWNAME=). If necessary, when the copy is complete, the original cluster can be deleted and the new cluster (and all of its components) can be renamed with an IDCAMS ALTER; however, it is easier to dump and restore the cluster with FDRDSF. This is an example of copying a 2-volume cluster, using CATDSN= to select the input disk volumes and using NVOL= to specify the output:

```
//COPY          EXEC   PGM=FDRCOPY,REGION=2M
//SYSPRINT      DD     SYSOUT=*
      COPY TYPE=DSF
      S CATDSN=MULTI.VOLUME.CLUSTER,NEWI=.+NEW,NVOL=(TEST01,TEST02)
```

For both FDRDSF and FDRCOPY, the volumes of the multi-volume cluster can be processed in one step or in multiple steps, and in any order. But the cluster will not be usable until all volumes have been processed.

For restore from any type of ABR backup, simply provide one SELECT statement specifying the base cluster name, e.g.,

```
SELECT DSN=cluster
```

or a generic selection mask which selects the clusters desired.

If the cluster still exists on disk, ABR will locate all volumes it exists on, and restore the contents of each component on all volumes. If the cluster does not exist, ABR will default to restoring to its original volumes unless you override with the NVOL= operand.

Multi-volume clusters cannot be restored or copied/moved to an unlike device type, e.g., 3390 clusters can only be output to 3390s. FDRREORG, IDCAMS or other utilities must be used if you must move multi-volume clusters to new device types.

80.14 GENERIC DATA SET NAME SELECTION

On many FDR control statements, a function is available which allows you to select data sets for processing using a generic name, also called a "data set name filter". The documentation for various FDR statements will refer to this section wherever it is supported. Do not assume it works for every statement, check the documentation.

Generic data set name selection allows you to specify a filter which will be applied to all data sets selected from system catalogs, VTOCs, backup files, or the Archive Control File. The filter allows you a great deal of flexibility in specifying the names to be selected.

FILTER CHARACTERS

The data set name filter is a string of characters which specify the data sets to be selected. Each character in the filter may represent a single character in the name of the data sets or may represent a variable number of characters. The filter characters may be:

- any alphanumeric character (A-Z or 0-9) or national character (\$ # @ in the US) represents itself and must appear in the indicated position in the data set name
- / (slash) or % (percent) represents ANY single valid alphanumeric or national character
- | (vertical bar) represents any single alphabetic (A-Z) character
- + (plus) represents any single numeric (0-9) character
- ? (question) represents any single national (\$ # @ in the US) character
- . (period) is used to separate index levels. For compatibility with earlier releases, if the filter begins with one or more periods, this indicates that that many index levels at the beginning of the name are to be skipped over before applying the rest of the filter
- * (asterisk) by itself as an index level indicates that the index level must exist in the selected name, but that it can contain any valid characters and can be any valid length (1 to 8 characters). However, if a single asterisk is combined with other characters in an index level, then it represents a variable length string (zero or more characters) at that point in the index level
- ** (double asterisk with no trailing period) represents zero or more characters or index levels. The filter routine will try to apply the remainder of the filter to the data set name beginning with the current character and stepping through the name until it matches or until the end of the name is reached
- **. (double asterisk with a trailing period) also represents zero or more characters or index levels, but the filter routine will try to apply the remainder of the filter to the data set name at the beginning of each index level (if not currently at the beginning of a level, it will start at the next level).

These special cases apply to ** with periods:

- **. at the beginning of the filter represents zero or more whole index levels at the beginning of the name
- **. at the end of the filter represents zero or more whole index levels at the end of the name.

Note that in the XDSN= operand of FDREPORT, the implementation is slightly different: "***." and ".***" represent ONE or more index levels at the beginning or end of the name.

The examples that follow will make this clearer, and will show you how flexible filtering can be.

80.14 CONTINUED

The examples below are designed to illustrate the power and flexibility you have in using FDR's generic data set name selection.

EXAMPLE 1

ACCOUNTS.PAY**

will select data set beginning with "ACCOUNTS.PAY" such as

ACCOUNTS.PAY.OVERDUE
ACCOUNTS.PAYROLL.CHECKS

EXAMPLE 2

USER1.*.CNTL

will select any 3-level data set name with USER1 as the first level and CNTL as the third, such as

USER1.JCL.CNTL
USER1.TEST.CNTL

EXAMPLE 3

USER+.**.*LIST

will select any data set whose first index is USERn (n is numeric) and whose last index level ends in LIST (including LIST by itself) with any number of index levels between them, such as

USER1.LIST
USER3.ISPF.OUTLIST
USER9.TEST.PRINT.MYLIST

EXAMPLE 4

TEST

will select any data set with the string TEST anywhere in the name, such as

ABC.TEST.GHI.XYZ
TEST2.LIST
REPORT.CURRENT.TEST
USER3.MYTEST

EXAMPLE 5

A?*B**

will select any data set whose name contains an A, a national character, and a B, in that order, such as

SYS1.A.X#B
APPLIC1.A@B.LOAD
APPLIC2.MASTER.BALANCE\$.BOUNCE

EXAMPLE 6

PAYROLL.%%||+*.FILE

will select any three-level data set whose first index is PAYROLL, whose second index consists of any 2 characters, followed by any 2 alphabetic characters, followed by any 1 numeric character, and up to 3 more characters, and whose third index is FILE, such as

PAYROLL.DCLX5.FILE
PAYROLL.A4TV3LM.FILE

EXAMPLE 7

PAYROLL.//***.FILE.**

will select any data set whose first index is PAYROLL, whose second index consists of at least 2 characters, and which has any number of following index levels, one of which must be FILE, such as

PAYROLL.DCLX5.FILE
PAYROLL.AB.FACTORY1.FILE.OCT90
PAYROLL.FACTORY2.HISTORY.FILE

With the same mask, FDREPORT will select only data sets with at least one index level after the index FILE.

80.14 CONTINUED

EXAMPLE 8 ** . G++++V00

will select all generations of any GDG (Generation Data Group). In ABR, you can do this more easily with the GDG operand, but this may be useful in other programs which support generic data set selection.

EXAMPLE 9 **MASTER (- 1)

will select the -1 generation of any GDG whose name ends in MASTER, such as

```
PROD1.WEEKLY.WIDGET.MASTER ( - 1 )
FINANCE.GLEDGER.XMASTER ( - 1 )
```

This will only work when selecting data sets from the catalog (such as the CATDSN= operand).

EXAMPLE 10 **

will select all data sets.

WARNING: Use ** with care, especially when doing ARCHIVE or SUPERSCRATCH.

80.14 CONTINUED

**SELECTING
DATA SETS
FROM THE
CATALOG**

When generic data set selection is used to select data sets from system catalogs (such as the CATDSN= operand in ABR and FDRCOPY or the DATATYPE=CATALOG or CATVTOC in FDREPORT), one or more system catalogs will be searched depending on the filter specified.

By default, the search will start with the master catalog; if aliases are encountered which match the filter, the associated user catalogs will be searched as well. If there are enough explicit characters at the beginning of the filter, the search may be limited to one or a small number of user catalogs. For example, a filter of AB** means that only those user catalogs associated with aliases beginning with AB need be searched. However, if the filter begins with generic characters (such as **AB), then every user catalog with an associated alias in the master catalog must be searched. Although the catalog search is designed to be as efficient as possible, reading the catalogs directly whenever possible, this may still be a time-consuming process.

You can specify the name of a catalog to search instead of going through the master catalog; in ABR and FDRCOPY specify the CATALOG= parameter. In this case only data sets in that catalog will be selected; no other catalogs will be searched even if they exist as user catalogs in that catalog. If you want to start your search with a master catalog which is not the master catalog of the system you are executing on, specify MCATALOG= instead and aliases will be searched.

GDGs (generation data groups) are treated specially when selected from a catalog:

- 1) the filter will be compared to the GDG base name (without the GnnnnVnn absolute generation). If it matches the base, then all generations in that GDG will be selected (unless limited by CATLIMITGDG= as shown below).
- 2) the filter will also be compared to the full name of each generation (including the GnnnnVnn). All generations which match the filter will be selected.
- 3) if the filter ends in a relative generation number in parenthesis (e.g., "PAYROLL.**(-n)" or "PAYROLL.**(0)") then only that generation will be selected from any GDGs selected (and the filter will not select any non-GDGs).
- 4) In ABR and FDRCOPY you can specify CATLIMITGDG=-n and only the (-n) generation of any GDGs which match the filter will be selected. However, the filter may also select non-GDG data sets.
- 5) In ABR and FDRCOPY you can also specify CATLIMITGDG=n and only the most recently created "n" generations of any GDGs which match the filter will be selected. The filter may also select non-GDG data sets.
- 6) For SMS-managed GDGs, any generations in rolled-off or deferred roll-in status will NOT be considered part of the GDG for selection; they will be treated as normal non-GDG data sets.

For ICF VSAM clusters, the filter will only be compared to the base cluster name. Component names, PATH names and AIX (alternate index) cluster and component names will not be tested.

It is possible to assign aliases to non-VSAM data sets. These aliases are complete alternate names by which a data set can be referenced and are not related to the aliases assigned to catalogs to control which data sets will be in a given catalog. The filter will never select alias names; it is tested only against the true name of every data set.

Filter processing for catalogs supports ICF VSAM catalogs. It also supports old non-ICF catalogs for non-VSAM data sets only. OS CVOL catalogs (SYSCTLG) will be bypassed, with the exception that if the filter specifies an explicit data set name, with no special selection characters, it will select that one name even if it is in a CVOL; this also supports selecting a single GDG generation from a CVOL by relative generation number.

STPCAT and JOBCAT, if present, will be normally be ignored. However, if the mask begins with a complete unqualified index, plus at least 1 more character after the period (e.g., CATDSN=PROD.X**), and the STPCAT/JOBCAT contains data sets which match the high-level index, then only the STPCAT or JOBCAT will be searched for that request.

80.14 CONTINUED

If an alias is encountered in a master catalog which points to a user catalog on a volume which is not currently mounted (or does not exist) that catalog will be bypassed.

By default, any errors scanning the catalogs (errors opening a catalog, internal errors in a catalog, or an alias which points to a CVOL) will result in a diagnostic message. In ABR and FDRCOPY, the operand CATBYPERR will bypass printing errors for CVOLs and catalog OPENS; internal errors will still be reported.

**SELECTING
DATA SETS
FROM OTHER
SOURCES**

When generic data set selection is used to select data sets from sources other than the system catalogs (such as VTOCs or the Archive Control File), then the entire source will be scanned, since the names may be in random order. Each data set name taken from the input will be processed against all of the data set filters (such as multiple SELECT statements in ABR, DSF, and FDRCOPY).

GDG generations will be treated as normal data sets. There is no provision for selecting by relative generation number. It is not sufficient for the filter to match on just the GDG base name; the filter must take into account the GnnnnVnn absolute generation number at the end (e.g., "PAYROLL.MASTER.**" or "PAYROLL.MASTER.G*").

For ICF VSAM clusters, the filter will only be compared to the base cluster name; component names and AIX names will not be tested. However, FDREPORT will test component names.

80.15 SECURITY

WARNING: by default no security checks are done for FDR operations, with the exception of a few checks done by operating system components. In general there is no security for FDR operations unless you enable FDR security checking via the ALLCALL option in the FDR Global Option Table as described in [Section 90.12 "Security Options"](#).

This section details the security checks that are enabled by the ALLCALL option. In addition, many other sections of this manual contain a summary of security checks done by the operation described in that section.

All of the security checks described below are issued via the RACROUTE macro. RACROUTE issues a call to an MVS component called SAF (Systems Authorization Facility), which routes the call to the active security system. This allows FDR to do security checking without needing to know which security system is involved. SAF is supported by RACF, CA-TOP SECRET, CA-ACF2 and others.

DASDVOL VOLUME PROTECTION

For most functions, FDR checks for a volume profile (CLASS='DASDVOL',ENTITY='volser') to see if the user is authorized to the entire volume; please read the following descriptions of security checks by FDR function for details.

Innovation recommends that you establish DASDVOL volume profiles for all of your DASD volumes, and authorize the userids that will be performing volume level operations (such as backups, COMPAKTions and full-volume restores) with appropriate access to those profiles for 2 reasons:

- In most cases, users must have appropriate authority to every data set being processed by FDR. When a large number of data sets are selected (such as a full-volume backup) the overhead to security check every one can be considerable. But, if the user is authorized to the entire volume under a DASDVOL profile, it is assumed that they have the same authority to every data set on the volume, so individual data set security checks are bypassed, greatly reducing that overhead. In these cases, since a failure of the DASDVOL check simply leads to individual data set checks, the DASDVOL check is issued with a parameter (LOG=NOFAIL) which suppresses any error messages or logging associated with it, since this could erroneously be interpreted as a security violation.
- In some cases, the **only** security check done is a DASDVOL check, so for volumes that are not protected by a DASDVOL profile, any user will be able to do the operation. These DASDVOL checks are done without LOG=NOFAIL, so potential security violations are logged.

RACF users who use the GDASDVOL resource grouping class must insure that this option is in effect:

```
SETROPTS RACLIST(DASDVOL)
```

For CA-TOP SECRET and CA-ACF2 users, consult vendor documentation for instructions on implementing DASDVOL class profiles.

80.15 CONTINUED

**DUMP
TYPE=FDR**

For full-volume backups under FDR or ABR, FDR will first check for READ authority to the input disk volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
ATTR=READ
```

If a DASDVOL profile does not exist for the volume, or the user is not authorized under the DASDVOL profile, FDR will then check for READ authority to every data set on the volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
VOL=volser,ATTR=READ
```

If the user is not authorized to read **every data set** on the volume (either individually or globally via DASDVOL), the backup will be terminated. Since full-volume backups always backup every data set on the volume, there is no facility to skip unauthorized data sets.

**RESTORE
TYPE=FDR**

For full-volume restores under FDR or ABR, FDR will check for ALTER authority to the output disk volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
ATTR=ALTER
```

If a DASDVOL profile exists for that volume, and the user is not authorized to alter the volume, the restore will be terminated. No individual data set checking is performed. **Note that this means that without DASDVOL protection, any user can do a full-volume restore.**

The user's authority to read the disk volume that was backed up is checked by OPENING the backup file (see "Additional Security Considerations" later in this section).

**COPY
TYPE=FDR**

For full-volume copies, FDR will perform the same checks as DUMP TYPE=FDR on the input volume and the same checks as RESTORE TYPE=FDR on the output volume.

COMPAKTOR

For COMPAKT TYPE=FASTCPK and TYPE=RELEASE, COMPAKTOR will check for ALTER authority to the volume being processed with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
ATTR=ALTER
```

If a DASDVOL profile exists for that volume, and the user is not authorized to alter the volume, the operation will be terminated. No individual data set checking is performed. **Note that this means that without DASDVOL protection, any user can do a FASTCPK or RELEASE.**

For an old-style COMPAKTion from a backup data set (COMPAKT TYPE=CPK, which is the default) COMPAKTOR does the same DASDVOL check on the output volume as described above. For a DUMP=NO COMPAKTion (from an existing backup) this is again the only check done. But if an FDR backup is taken as part of the CPK run (the DUMP=YES operand to invoke a FDR backup, or the COMPAKT option on an FDR or ABR full-volume backup), the FDR backup will perform the checks described above for DUMP TYPE=FDR on the input volume (usually the same as the output volume); this will include individual data set checks for READ authority if the user does not have DASDVOL authority to the entire input volume.

80.15 CONTINUED

**DSF ABSOLUTE
TRACK
OPERATIONS**

When the FROM/TO operands are used in FDRDSF or FDRCOPY to access tracks by their absolute address, FDR will check for authority to the volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
        ATTR=access
```

where "**access**" is READ for DUMP and PRINT and for the input volume on COPY and is ALTER for RESTORE and for the output volume on COPY. If a DASDVOL profile exists for the volume being accessed, and the user does not have the requested access authority, the operation will be terminated. FDR does not attempt to determine what data sets the requested tracks belong to, or to perform checking at the data set level.

If the installation decides that this approach does not provide adequate security, then absolute track operations can be disabled completely, as discussed under NOABSTRK, in [section 90.12](#).

**DUMP
TYPE=DSF/ABR
/AUTO/APPL**

For all DSF and ABR data set backups (including application backups), FDR will first check for READ authority to the input disk volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
        ATTR=READ
```

If the user is authorized, the dump proceeds with no further checking. If a DASDVOL profile does not exist for the volume, or the user is not authorized under the DASDVOL profile, FDR will then check for READ authority to each data set which is selected for backup with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=READ
```

Any data sets to which the user is not authorized will be bypassed (not backed up). All authorized data sets and all unprotected data sets will be processed.

**DUMP
TYPE=ARC/SCR**

For ABR ARCHIVE and SUPERSCRATCH, FDR will do the same checks as for DUMP TYPE=DSF described above, except that the requested access authority is ALTER since the data set will be scratched.

However, DUMP TYPE=ARC,SCRATCH=NO is an application backup, equivalent to DUMP TYPE=APPL, so the requested access authority is READ.

80.15 CONTINUED

**RESTORE
TYPE=DSF/
ABR/ARC/
APPL TO SAME
NAME**

For all DSF and ABR data set restores, FDR will first check for UPDATE authority to each output disk volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
ATTR=UPDATE
```

If a data set being restored **does not exist** on disk before the restore, then FDR will request the Operating System to create the data set. The Operating System will check for CREATE authority at the data set level, whether or not the user has DASDVOL authority, and regardless of what security options the installation has activated within FDR. If the user is authorized to create the data set, FDR assumes they must be authorized to write to it; therefore FDR will not issue its own security check.

If a data set being restored **does exist** on disk before the restore, and a DASDVOL profile does not exist for the volume, or the user is not authorized under the DASDVOL profile, then FDR checks for UPDATE authority to the data set with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
VOL=volser,ATTR=UPDATE
```

If a particular data set fails the security checks, that data set will be bypassed, but other requested data sets may still be restored.

If the restore is an automatic recall of an archived data set, then additional options are available to modify the security checks. For example, you can allow the restore even though the user has only READ authority so that data sets can be recalled and read, or you can allow the restore even though the user has no authority to the data set. For details, [see Section 51.35](#).

**RESTORE
TYPE=DSF/
ABR/ARC/
APPL TO NEW
NAME**

When restoring data sets to a new name, FDR will check authority to both the original data set (to verify authority to read that data set) and output dataset (to create or update it). READ authority to the original data set name on the original volume serial is checked with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
VOL=volser,ATTR=READ
```

If the user is not authorized to the original name or the data set is not protected on the current system, FDR will **not restore** the data set. Other data sets requested in the same run will still be restored.

The reason why FDR fails the restore if the original dsname is unprotected (return code 4 from RACROUTE indicating no profile exists) is that it is possible that the data set was protected by a profile at the time it was backed up, and that the user running the restore would not have been authorized under that profile, but the profile has since been deleted or the restore is running on a different system. At the time of the restore, FDR cannot tell whether the data set was protected at the time of the backup. So if the input data set is not *currently* protected, the restore is disallowed. However, this means that if a data set has never been protected, FDR will not allow anyone to restore it to a new name; if you have unprotected data sets in your installation, contact Innovation for assistance.

No checking is done for a DASDVOL profile for the volume from which the data set was backed up; so if the user has authority under a DASDVOL profile for the original volume, he must still have to have authority under a data set profile for the original dsname in order to restore the data set to a new name.

If the user is authorized to read the original data set name, FDR will check security for the output volume and the newname output data set as described above for RESTORE TYPE=DSF TO SAME NAME.

80.15 CONTINUED

ABR AUTO-RECALL For auto-recall, there are additional security options in addition to those described above. See [Section 51.35](#).

COPY/MOVE FDRCOPY is essentially a dump and a restore in one task, so its security checks are essentially those described above for data set backups and restores.

For an FDRCOPY COPY operation, the input volume and data sets will be checked as described above for DUMP TYPE=DSF. The output volume and data sets will be checked as described above for RESTORE TYPE=DSF.

For an FDRCOPY MOVE operation, since the input data set will be scratched, FDRCOPY will first check for ALTER authority to DASDVOL on the input volume; if the user does not have that authority, it will check for ALTER authority to the data set itself. This is the same as the checks described above for DUMP TYPE=DSF except that ATTR=ALTER. The output volume and data sets will be checked as described above for RESTORE TYPE=DSF.

On a COPY or MOVE to a new name, the input data set is checked according to the rules just described. The special rules described above for RESTORE TYPE=DSF TO NEW NAME do not apply.

FDRABRM FDRABRM is the ABR VTOC maintenance utility, which sets disk volume processing options for ABR. Under ISPF, it is invoked by option A.I.8. FDRABRM will first check for ALTER authority to the disk volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
        ATTR=ALTER,LOG=NOFAIL
```

If a DASDVOL profile exists and the user is not authorized under the DASDVOL profile, FDRABRM will terminate processing of that volume. No security violation message will appear on the job log, but an error message will appear in the SYSPRINT listing.

If the user is authorized under the DASDVOL profile or there is no DASDVOL profile, FDRABRM will then check for ALTER authority on the ABR model DSCB with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=ALTER
```

FDRABRM will then check for READ authority to every data set on the volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=READ
```

If the user is not authorized to read **every data set** on the volume (either individually or globally via DASDVOL), the backup will be terminated.

ABR REMOTE QUEUE OPERATIONS In ABR, requests for backups, archives, and restores can be added to "remote queues" by end users. Later, those remote queue requests are processed by a job submitted by the data center or some other central group. Requests are added by the ABR Remote Queue Utility (FDRABRUT), which can be invoked directly by a user, by an ISPF dialog, or by an ABR auto-recall for a TSO user.

The security checks described below are done under the userid of the requesting user, to verify that the user has the required authority to the data sets they are requesting. Note that users may also need to be authorized to update the Remote Queue data sets (see "Additional Security Considerations" below).

The jobs which process the remote queues will usually be run under a userid associated with Operations. Since the normal checks for backups, restores, and ARCHIVES described earlier will be done, that job must be authorized to all volumes or data sets that the users have added to the remote queues.

80.15 CONTINUED

**ABR REMOTE
QUEUE
BACKUP or
ARCHIVE**

If DISKUPDATE=YES is in effect (see [Section 90.15](#)), then the Remote Queue Utility (FDRABRUT) or the ABR ISPF dialog, before setting the Update indicator or ARCHIVE indicator in the DSCB, will check for authority to the volume with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DASDVOL',ENTITY=volser,
        ATTR=access
```

where "access" is READ if all of the commands for a given volume are BACKUP or RESET BACKUP, or ALTER if at least one of the commands for the volume is ARCHIVE or RESET ARCHIVE. If a DASDVOL profile does not exist for the volume, or the user is not authorized under the DASDVOL profile, then authority to each data set that is selected will be checked with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=access
```

where "access" is READ if the command for that data set is BACKUP or RESET BACKUP, or ALTER if the command is ARCHIVE or RESET ARCHIVE.

If DISKUPDATE=NO is in effect, then the Remote Queue Utility or the ABR ISPF dialog, before adding the request to a Remote Queue data set, will not check for volume authority, but will check authority to each data set that is requested to be processed with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=access
```

where "access" is READ if the command for that data set is BACKUP or RESET BACKUP, or ALTER if the command is ARCHIVE or RESET ARCHIVE. If the data set is not cataloged and VOL is not specified by the user, then the request will fail since the volume serial is unknown. If DISKUPDATE=NO is in effect and ALLCALL is enabled then only requests for fully-qualified data set names are allowed because proper security checking cannot be done if DSG= or DSN=mask is used; VOLG= is also disallowed for the same reason.

If a particular data set fails the security checks, that data set will be bypassed, and other requested data sets may still be processed.

**ABR REMOTE
QUEUE
RESTORE**

Remote Queue RESTORE requests are always processed by adding the request to a Remote Queue data set. The Remote Queue Utility or the ABR ISPF dialog, before adding the request will check data sets with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
        VOL=volser,ATTR=access
```

It will check the original dsname on the original volume serial with "access" of READ.

If the user specified a new output data set name or a new volume serial, RACROUTE will be also issued for the new data set with "access" of ALTER. If a profile exists for the new data set, then FDR will require that a profile exist for the original data set also. If the original data set is not cataloged and VOL is not specified, or if the new data set is not cataloged and NVOL is not specified, then the request is failed. If ALLCALL is enabled then only requests for fully-qualified data set names are allowed because proper security checking cannot be done if DSG= or DSN=mask is used; VOLG= is also disallowed for the same reason.

If a particular data set fails the security checks, that data set will be bypassed, and other requested data sets will still be processed.

80.15 CONTINUED

FDRREORG For IAM and VSAM files that it processes, FDRREORG will check for authority with:

```
RACROUTE REQUEST=AUTH,APPL='FDR',CLASS='DATASET',ENTITY=dsname,
VOL=volser,ATTR=ALTER
```

A REORG of a PDS with FDRREORG (or an FDRCOPY "REORG TYPE=DSF") requires UPDATE authority to the volume or the PDS, the same as defined for "Restore to Same Name" above.

In addition, FDRREORG uses dynamic allocation and OPEN:

to create and open its backup and log data sets, using names defined by FDRREORG parameters and defaults ([see Section 90.23](#)). It must have CREATE authority to those data sets.

for the VVDS on every volume containing VSAM being reorganized. It requires READ authority to the SYS1.VVDS.Vvolser data sets.

These checks are done by the operating system, so these authorizations are required even if the ALLCALL option is not in effect.

**ADDITIONAL
SECURITY
CONSIDERA-
TIONS**

Normal OPENs are done on backup data sets, so OPEN will check that the user has authority to create the backup for DUMPs and to read the backup on RESTOREs. However, security checking for tapes and tape data sets may be optional so check your security system documentation for details on tape security. You may wish to restrict most users from reading backup data sets by protecting the backup data set names.

When the ABR DADSM Pre-processing exit is installed, it bypasses RACF and Top Secret security checking when it catalogs entries into the ABR SCRATCH catalog. ACF2 users must either install the ACF2 VLDEXIT pre-validation exit ABRVALD (supplied by Computer Associates, and documented in the ACF2 Other Products manual), or be authorized to create entries in the ABR SCRATCH catalog, i.e. they must be authorized to catalog dsnames starting with '#.'.

For all of the security checking described in this section, FDR supports both "discrete" and "generic" RACF security profiles. Although generic profiles are used in most RACF installations, discrete profiles (one per data set) may still be in use. Discrete profiles may be automatically deleted when the data set is deleted, and FDR does **not** preserve discrete profile information on its backups, so when a data set that was protected by a discrete profile is restored by FDR, a "default" discrete profile will be created for it, granting access only to the user doing the restore; this profile may require manual updating.

Innovation recommends that the profile or security rule protecting a data set or volume should not be deleted as long as any copy of that data set or volume exists on a backup or ARCHIVE tape.

FDR jobs which use HFS=QUIESCE to quiesce HFS file systems must be authorized to use OS/390 Unix services. For details, see "HFS" in [Section 80.11](#).

80.15 CONTINUED

**RESTRICTING
PROGRAM
ACCESS**

RACF has a feature called PADS ("program access to data sets") which can restrict the program names which can open particular data sets; other security systems have a similar feature (sometimes called "program pathing"). If you want to limit access to certain data sets to FDR programs, you will need to define rules which allow only programs starting with FDR to open them. PADS requires that every program loaded at the time the file is opened be explicitly listed; if you wish to use PADS to restrict access to certain data sets to FDR programs, you can define every program in the FDR program library which starts with FDR (plus IGG019YZ), or contact Innovation for a list of modules used for specific functions. Other security systems may be able to conveniently define rules for generic FDR* programs; check your security documentation.

FDR and ABR backups contain the data from many data sets. You may need to grant users READ authority to those backups in order to restore data sets to which they are authorized, but you don't want them to be able to open or browse the backup data so that they cannot see nonauthorized data sets. So, you may wish to restrict access to backup data sets to FDR programs. You can eliminate the need to give READ authority to users for ABR backups by setting options so that all recalls and restores are done in the background or via Remote Queues.

Users who are allowed to restore data sets from ARCHIVE will need UPDATE authority to the ARCHIVE Control File (to set a "restored" flag), but you may want to restrict them to reading or updating the Control File with FDR programs.

Users who will add requests to the ABR Remote Queue data sets require UPDATE authority to them, but you may want to restrict them to reading or updating the Remote Queues with FDR programs. It may also be possible to bypass security checking during the Remote Queue updates; contact Innovation for details.

80.20 DISASTER RECOVERY

One of the purposes of data backup with a product such as FDR is recovery of the data necessary to keep your business running at a disaster/recovery site after a disaster of some sort disables your home Data Center.

RECOVERY PROCEDURES

This section details several procedures for recovering data from ABR backups after a disaster. These procedures may be used to recover all or a critical subset of your data at a disaster recovery site, and may also be used when you move back to your restored home site. These procedures may also be useful when you are simply moving to a new data center location.

- 1) SAR (Stand-Alone Restore) may be required to restore full-volume backups of the volume necessary to IPL your operating system. This may include the IPL volume (sysres), spool volumes, paging volumes, catalog volumes and others. If your disaster/recovery site provides a "starter" system, you can execute FDR or ABR on that system to do the same restores.
- 2) ABR full-volume recovery recreates entire disk volumes on disks at the new site. In each case the recovered volume will look exactly like it did at the time of the most recent incremental backup. This is the most efficient way of rebuilding your data. However, it does require that the new site have a disk volume of the proper type (e.g., 3390) and size (or larger) for each disk you plan to recover.
- 3) ABR data set restore is used to restore the latest copy of individual data sets from the full and incremental ABR backups. This procedure allows you to restore selected data sets to new volumes, even if the new volumes are smaller or larger than the original volumes (as long as individual data sets will fit on the selected volumes). This can even be used to do "logical" restores to a different device type (such as 3380 to 3390) if the original device type is not available; however, this is not recommended since it is less efficient and has restrictions as documented in [Sections 80.12](#) and [80.13](#).
- 4) ABR Application restore is used to restore the latest copy of individual data sets which were backed up by ABR Application Backup ([see Section 52](#)). Since Application Backups are run as part of an application jobstream, the data sets will be restored to a point-in-time that is most useful to that application.

FDRCLONE

FDRCLONE is an optional, extra-cost feature of FDR which simplifies many of the procedures described in this section. It has 2 major functions:

- the FDRCLONE program can be used to dynamically restore application data sets from ABR volume backup tapes as they are needed, instead of doing full-volume recoveries which probably restore many data sets which will **not** be used. It operates using techniques similar to ABR auto-recall; when a batch job or TSO user references a "cloned" data set which has not yet been restored to disk, ABR is invoked to restore it from the appropriate backup tape. This allows you to start work as soon as your required system volumes are restored. FDRCLONE may be able to fit the required application data sets on fewer volumes than they occupied at your home site, since only referenced data sets are restored.
- FDRCLONE also includes the program FDRDRP (Disaster Recovery Program). FDRDRP does full-volume recovery from ABR full-volume and incremental Volume Backups, just like ABR does (see above). While ABR does one disk volume at a time in a given ABR job and may mount the required tapes many times, FDRDRP does many disk volumes in parallel and mounts the tapes a minimum number of times, restoring data belonging to many disk volumes in one pass of the tape. This can greatly reduce the time required to do full-volume restores.

FDRCLONE is described in starting in [Section 50.70](#).

80.20 CONTINUED

**CHOOSING A
PROCEDURE**

You may need to use several or perhaps all of these recovery procedures to rebuild your Data Center at the disaster site:

- 1) If your disaster site does not have a starter MVS system, you must use SAR to restore the full-volume backups of the minimum volumes required to IPL your system. If you do have a starter system available, you can do FDR full-volume restores of those same volumes, and then IPL your system to complete the recovery (you could also restore all your ABR data on the starter system).

WARNING: as of OS/390 V2R4, your system may not IPL unless a root HFS (Unix file system) data set is present. Be sure to include the root HFS volume among those to be restored.

- 2) If your disaster site has sufficient disk volumes of the proper type and size, you will probably want to do ABR or FDRDRP full-volume recoveries for each disk volume. This is the most efficient method of restoring those volumes to their exact contents as of the time of their last incremental backup. If the disaster site has disks of different types or smaller disks or a fewer number of larger disks, you may not be able to use full-volume recovery for your entire installation.
- 3) If your disaster site has less capacity than your home site, or a smaller number of larger disks, you will need to do ABR data set recovery or FDRCLONE so that the restored data sets can be allocated on those disk volumes. You can also selectively restore critical data sets if the disaster site does not have room for all of your data.
- 4) For applications which have implemented Application Backup, you will probably want to use Application restore, since it will restore the data sets to the point-in-time designated by the application (data sets restored with ABR full-volume or data set restore will be restored to the time of the last incremental backup, chosen by the Data Center).

80.20 CONTINUED

**REQUIRED
RESOURCES**

The disaster recovery procedures assume that several resources were stored in some offsite storage location and are available at the disaster/recovery site. Please review this list carefully to be sure that you will not be missing some required item:

- 1) Copies of all of your ABR Volume Backups (full-volume and incremental backups). These are usually the COPY2 backups produced by ABR using the TAPExx DD statements, but they may also be copies (COPY2 through COPY9) produced by the FDRTCOPY or FDRTSEL utilities ([See Section 60](#)). At least the most recent full-volume backup of each required volume must be in the offsite storage, plus as many of the subsequent incremental backups of those volumes as possible.
- 2) A backup of a system residence volume (SYSRES) which is configured to run on the hardware of the disaster/recovery site. You must generate an I/O configuration, with HCD or IOCP/MVSCP, which matches the configuration of the site. It is possible to put multiple configurations on the same SYSRES, so this might actually be your normal SYSRES volume. In some cases, multiple volumes are necessary to contain the configurations.
- 3) A separate backup of the ABR catalog produced by FDRDSF or by ABR DUMP TYPE=DSF, after all daily ABR runs are complete (for safety, send 2 copies of this backup offsite). Do not depend on the backup of the ABR catalog made by normal ABR incremental backups, since this will be taken before all ABR backups have been recorded.
- 4) A hardcopy of the ABR catalog, produced by program FDRABRP, using the command PRINT CATLG ([See Section 53](#)). This will serve as the ultimate backup if the ABR catalog backups are unusable, and can help you prestage the required tapes. If you must do SAR restores it will identify the volumes containing the latest full-volume backups.
- 5) A tape containing SAR (the Stand-Alone-Restore program of FDR). It will be easiest to IPL SAR if the SAR tape is an unlabeled tape, but a labeled tape can also be used, or a copy of the FDR distribution tape from Innovation which contains SAR as file 1. [See Section 15](#) for instructions on generating a SAR IPL tape.
- 6) A tape containing an IEBCOPY unloaded copy of the ABR program library may be required if you are going to use an operating system provided by the disaster recovery site to do the restores, since they may not have a current copy of FDR or ABR available.
- 7) The tape containing the daily FDREPORT extract file, if you are using the data set recovery procedure described later in this section. If you are using FDRCLONE, this will be the tape containing the FDRCLONE data file.
- 8) Separate backups of other special-purpose data, such as tape management system data bases, using special utilities, may be required to successfully restore data with ABR; however, we suggest that you do not activate tape management or security systems until the ABR restores are done.
- 9) ABR Application Backup tapes for all applications using ABR Application Backup.

The more frequently these resources are sent offsite, the more current the offsite information will be, and so the more current your system will be after restoring at the disaster/recovery site. The expense and effort of sending data into offsite storage must be evaluated by each installation versus the impact of having back-level data after disaster recovery.

80.20 CONTINUED

**FULL VOLUME
RECOVERY**

The following procedure may be used to recreate entire DASD volumes from the ABR Volume Backups. You can use this to recover your entire installation, to recover a critical subset of your volumes, or just to recover the volumes required to operate your system (the latter assumes that you will recover critical data by other means, such as the ABR data set restore or ABR Application Restore).

STEP 1 This creates an IPLable version of your operating system. You can skip this step if the disaster site provides you with a starter MVS system adequate for ABR restores.

Do Stand Alone Restore (SAR) of the system residence volume, any other volumes that you need in order to IPL (such as paging and spool volumes, and volumes containing LPA and linklist libraries) and the volume containing the ABR program library. The input to each Stand Alone Restore is the most recent full-volume backup (TYPE=FDR) for that volume. This step will not necessarily restore the volumes to the most current level, but at this stage you are just trying to get a system you can IPL. Remember that the system residence volume you restore must be one that supports an I/O configuration that will run on the disaster/recovery site hardware.

STEP 2 IPL the operating system (either the system you just restored, or the starter system provided by the site).

If the system you have IPLed does not contain a current ABR program library, create that program library. You can reinstall the ABR library from the ABR distribution tape (following the instructions in [Section 90](#)) if you have an offsite copy of that tape, or you can restore an IEBCOPY unloaded copy of that library if you have sent that offsite. The ABR library must be APF authorized.

STEP 3 Next, the ABR catalog must be restored and made available to the system on which you will do the restores. Innovation recommends that you restore the ABR catalog to a spare volume, not the one on which it normally resides, so that it will not be overlaid by the full-volume restores in step 4.

If the master catalog of that system already contains the FDRABR alias or a catalog by the name of your catalog, it may be necessary to delete it with IDCAMS with this input:

```
DELETE (FDRABR,#) ALIAS
EXPORT catalogname DISCONNECT
```

Note: Your installation may have changed the FDRABR alias or the ABR scratch alias to another name; use the appropriate alias name.

Now the separate FDRDSF or ABR backup of the ABR catalog can be restored, to some DASD volume which does not currently contain such a catalog. DSF will allocate the catalog on that volume and do an IMPORT CONNECT for it in the master catalog of the running system. The aliases associated with the ABR catalog will be automatically defined. The JCL to do this restore will look something like this example:

```
//RESTCAT EXEC PGM=FDRDSF
//SYSPRINT DD SYSOUT=*
//TAPE1 DD DSN=catalog.backup.name,
// VOL=SER=serial,UNIT=TAPE,DISP=OLD
RESTORE TYPE=DSF
SELECT DSN=catalogname,NVOL=vvvvvv
```

Innovation recommends that you restore the ABR catalog to a volume other than the one on which it normally resides.

80.20 CONTINUED

STEP 4 Use the ABR full-volume restore function to reconstruct all of the required volumes at the most current level, using job steps similar to:

```
//RETVOL      EXEC   PGM=FDRABR
//SYSPRINT    DD     SYSOUT=*
//DISK1       DD     VOL=SER=nnnnnn,UNIT=SYSDA,DISP=SHR
//SYSIN       DD     *
      RESTORE      TYPE=FDR,CPYVOLID=YES,DYNTAPE2
      SELECT       VOL=vvvvvv,COPY=2,NVOL=nnnnnn
```

The ABR full-volume restore function will automatically recreate each volume exactly as it was at the time of the most recent incremental backup. It reads that most recent incremental backup first, restoring the VTOC and all other data on that backup, then reads backwards through the earlier incrementals until it reaches the full-volume backup that started the current generation, restoring only the data required from each backup. Unlike some competing products, in ABR there is no need to separately "apply" incrementals; the entire restore is done by the above step.

The output volumes do NOT have to be pre-initialized to the volume serials of your original volumes, nor do they have to be initialized for ABR. NVOL=nnnnnn specifies the current volume serial of the disk that will be overlaid by the restore of your volume "vvvvvv". CPYVOLID=YES will cause that volume to be renamed to your volume serial after the restore completes. Many disaster/recovery sites will guarantee that their disks will be initialized to a certain set of volume serials before you arrive; you can then plan that each one of your volumes (vvvvvv) is to be restored to a selected one of their serials (nnnnnn). The restore JCL must include a DISKxxxx DD specifying the output disk volser.

Change COPY=2 if COPY2 is not the copy you keep in your offsite storage. Note that you can change the default copy number for ABR restores to 2 so that you don't have to specify it in every job (see [Section 90](#), panel A.I.4.4).

The ABR model on each volume to be restored is not available (since the disks have not yet been restored), so ABR will locate the most recent incremental backup of the disk volume in the ABR catalog which you have just restored. ABR will then read that incremental backup, followed by each preceding incremental backup, and finally the full-volume backup which begins the current generation. Once this process is completed, the volume will look exactly as it did at the time that the latest incremental backup was taken. If there are certain volumes for which you want to restore only the most recent full-volume backup, add "GEN=CURRENT,CYCLE=00" onto the SELECT statement.

The DYNTAPE2 option specifies that ABR is to dynamically allocate 2 tape drives for restoring each disk volume, to improve performance. DYNTAPE2 is most useful when restoring from 3480/3490 tape drives, to avoid waiting for tapes to rewind. If you have 3490E or newer tape drives, or if you do not have enough tape drives to support the number of concurrent restores you want to do, change DYNTAPE2 to DYNTAPE, which will allocate only one drive.

If the backup tapes required for a given disk are in use by another restore job, ABR will wait for them to become available as long as you have the WTVOL option enabled in the ABR option table (See [Section 90](#), panel A.I.4.4). Use of DYNTAPE2 or DYNTAPE is recommended, since without it ABENDs may occur if the same tape is required by 2 concurrent restore jobs.

You can restore more than one volume in an ABR step (by specifying more than one SELECT statement), but those volumes will be restored one at a time, so it is easier to simply restore one volume per step as shown above. Those steps can be organized one per job, or several steps per job. In order to be able to restore a large number of volumes in minimum time, you will need to have multiple ABR restore jobs executing at once. Any number of concurrent ABR restores can be executing, limited only by the number of available tape drives. Performance may be reduced if the capacity of the tape or disk data paths is exceeded.

80.20 CONTINUED

However, since incremental backups will almost always put data from several disks on one tape, concurrent restore jobs will often contend for the same tape, causing one or more to wait. One way to avoid this is to divide the restores into groups of disk volumes, creating restore jobs which contain multiple ABR steps, each restoring one disk volume. The disk volumes in each restore group should be selected so that they will usually not be requesting the same backup tapes required by another restore group; this can be done in several ways:

- 1) the best way is to setup your normal ABR backup jobs to select disk volumes by the same groups. In other words, there will be a backup job which processes a given set of disk volumes, and a restore job which restores those same volumes, one at a time. Since the backup tapes used for one volume group will never overlap those used for another group, there will never be tape contention.
- 2) Another option would be to simply divide the volumes to be restored into groups based on the order in which they are selected for processing by your normal ABR backup jobs. For example, the first 10 volumes which ABR normally selects might be restored by one restore job, the next 10 by the next restore, etc. This will minimize contention but cannot totally eliminate it.

Note: if you use high-capacity tapes such as IBM 3590 Magstars or STK Redwoods, ABR may place the backups of many different disk volumes on one of those tapes. This may cause many of your restore jobs to serialize on that tape, causing the restores to run much longer than normal. In your DUMP jobs, you may wish to use the MAXFILE= operand to restrict the number of files that ABR will place on a tape to reduce this contention, but be aware that this will use more tape volumes and will waste more space on those volume.

FDRDRP ([Section 50.80](#)) can also be used to do these same full-volume restores. FDRDRP can restore many disks concurrently and mount the input tapes far fewer times than ABR. Although the operation for a single disk is the same under FDRDRP and ABR, FDRDRP attempts to mount each tape only once and restore the required data from all backups on the same tape. With FDRDRP, most of the preceding considerations do not apply. However, it may still be useful to do your backups in groups of volumes and do the FDRDRP restores in those same groups, to minimize tape contention.

STEP 5 If you used STEP 1 of this procedure, then you used SAR to restore the IPL volumes from ONLY the full-volume backup tapes, so some data sets on those volumes might be back-level. If this is a problem, then you must restore those volumes again with ABR. Since you are currently running your system from those volumes, the ABR restores must output to a different set of volumes; after the restores are complete, be sure that you place the original SAR-restored volumes offline and IPL using the new versions of those volumes (the IPL procedure will allow you to choose between two identically labeled volumes).

However, if your IPL volumes contain only data which does not change frequently, it may be unnecessary to do these repeat restores.

STEP 6 If you did not restore the ABR catalog to a temporary volume as recommended in STEP 3, and are using it on its normal volser, you must insure that the volume containing the ABR catalog is restored last. This restore cannot be run concurrently with the other restores because it may temporarily downlevel the catalog.

A downlevel copy of the ABR catalog was restored by the full-volume restores in STEP 4. When the full-volume restores are finished, repeat the restore of the ABR catalog that was done in STEP 3, to return the ABR catalog to the most current level. However, this time it should be restored to the volume where it normally resides. It is not necessary to delete the ABR catalog; the restore JCL in STEP 3 will simply overlay the existing catalog.

You have now completed the full-volume recovery procedure. If you are recovering all your data with full-volume recovery, you are done. If not, keep reading.

80.20 CONTINUED

**DATA SET
RECOVERY**

The following procedure may be used to restore individual data sets or groups of data sets from ABR full and incremental backup tapes. It can be used when you don't want to restore all of the data sets on those backups, or when you must restore data sets to volumes other than their original volumes because:

- the recovery site has volumes of different sizes, such as 3390-3s instead of your 3390-2s
- the recovery site has fewer volumes, requiring selective restores.
- the recovery site has disks of a different type, such as 3390s instead of your 3380s

Note that the last reason, which requires "unlike device" restores, has some limitations; please review [sections 80.12, 80.13](#), and member UNLIKE in the FDR ICL library for details. Unlike restore should be avoided if at all possible.

Data set restore at a recovery site requires that some additional information be saved as part of your daily backups. Normally, information about the backups of individual data sets is stored in the DSCB of the data set itself on disk. The ABR catalog contains the locations of the backup data sets created by ABR, but does not record information about the individual data sets in those backups. For disk volumes which have not been restored by the full-volume recovery process just described, ABR does not know where to find an individual data set.

However, if you tell ABR where to find a backup, specifying the disk volume, generation and cycle, e.g.

```
SELECT      DSN=dsname ,VOL=volser ,GEN=generation ,CYCLE=cycle
```

ABR can locate the backup and restore the data set.

FDREPORT, the ABR generalized report writer described in [Section 54](#), can build a "data base" which will contain the location of the latest backup of every data set in your system (or any subset of them). At the recovery site, you can then run FDREPORT against that data base to build SELECT statements as shown above to instruct ABR to restore selected data sets.

FDRCLONE ([Section 50.70](#)) is superior to the procedure described here for data set restores. FDRCLONE can automatically restore individual application data sets as they are referenced by batch jobs or TSO users.

80.20 CONTINUED

**BUILD THE
DATA BASE**

This jobstream should be run after all of your ABR full-volume or incremental backups are run each day, to build the data base which contains information about the current backup of data sets, but it should be run before the backup of the ABR catalog, since the data base file name starts with FDRABR and will be cataloged in the ABR catalog. The sample jobstream shown will gather information on every data set on every online disk volume in your installation. You can insert XSELECT and/or XEXCLUDE statements to limit the data sets it processes, or to limit the volumes it processes. You could also run several such jobs to build data bases about various subsets of your data sets. SYSUT2 is the database file, on tape so it can be sent offsite; it can be a GDG if you like.

```
//DATABASE      EXEC   PGM=FDREPORT,REGION=2M
//SYSPRINT       DD     SYSOUT=*
//ABRMAP         DD     SYSOUT=*
//SYSUT2         DD     DSN=FDRABR.OFFSITE.DATABASE,
//               UNIT=TAPE,DISP=(,CATLG),
//               DCB=(BLKSIZE=32760,BUFNO=10)
//SYSIN          DD     *
*      INSERT XSELECT/XEXCLUDE STATEMENTS HERE IF NEEDED
      REPORT   FIELD=(DSN,VOL,ABRGEN,ABRCYCLE)
      PRINT    ONLINE,RPTYPE=DATA,DISABLE=INFOMSG
*      GENERATE REPORT FROM EXTRACTED DATA
      DEFAULT  SORTALLOC=YES
      REPORT   FIELD=(DEFAULTS,ABRGEN,ABRCYCLE)
      SORT     FIELD=(DSN,VOL)
      PRINT    DATATYPE=EXTRACT
/*
```

Naturally, the data base tape(s) and the printed reports should be sent to your offsite storage with your other backups.

80.20 CONTINUED

**RESTORING
DATA SETS**

At the recovery site, FDREPORT is used to read the "data base" and generate SELECT statements for any data sets you wish to restore. These SELECTs are passed to an ABR restore step which will allocate the proper backup tapes and restore the data sets. The target volumes for the restores can be specified by NVOL parameters on the SELECTs (as shown in the example below) or by the ABR RESTORE ALLOCATE LIST (See [Section 90.32](#)). FDREPORT XSELECT statements can be used to select a subset of the data sets on the data base so that multiple restore jobs can be run concurrently.

The user catalogs in which these data sets are cataloged must be restored before this job stream is run. The RECAT and VRECAT operands will update those catalogs to point to the volumes to which the data sets were restored.

```
//SELREST      EXEC   PGM=FDREPORT,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSUT2       DD     DSN=FDRABR.OFFSITE.DATABASE,DISP=SHR
//SYSPUNCH     DD     DSN=&&ABR,UNIT=SYSDA,SPACE=(TRK,(5,2)),
//              DISP=(,PASS)
//SYSIN        DD     *
XSELECT XDSN=PAYROLL.**      ** select all PAYROLL data sets
PUNCH   FDLIB=MASK          ** punch mask for ABR statements
PRINT   DATATYPE=EXTRACT,RPTYPE=SELPCH
/*
//MASK         DD     *
)PREFIX
  RESTORE  TYPE=ABR,DYNTAPE,RECAT,VRECAT,COPY=2,MAXCARDS=2000
)ENDPREFIX
  SELECT  DSN=<NAME>,
          VOL=<VOL>,GEN=<ABRGEN>,CYCLE=<ABRCYCLE>,
          NVOL=PROD*
/*
//ABRREST      EXEC   PGM=FDRABR,REGION=2M
//SYSPRINT     DD     SYSOUT=*
//SYSIN        DD     DSN=&&ABR,DISP=(OLD,DELETE)
```

Even though you may be requesting many data sets from many different backup tapes, ABR will sort the requests so that each backup tape is mounted only once, if possible, and multiple data sets will be restored from a given backup file in one pass of that file, even if the outputs are going to multiple disk volumes.

Note that ABR always restores data sets to the same number of volumes they originally occupied, even if the output volumes are larger or smaller than the original volumes. For example, a 3-volume data set must be restored to 3 volumes, occupying the original amount of space on each volume. Likewise, a single-volume data set must be restored to a single volume and cannot be split across volumes. Restoration of large data sets will fail if they cannot be allocated on an appropriate number of volumes. If your installation has data sets which occupy entire 3390-3s, for example, they cannot be restored to a 3390-2.

80.20 CONTINUED**APPLICATION
RESTORE**

If some of the applications which run at your home Data Center have implemented Application Backup ([Section 52](#)) and send a copy of the backup tape offsite, then you will want to restore those backups. Application Backup allows the application programmers to choose the point in their daily processing where they want to take the backups; usually this is a point where it will be easy to restart the application. ABR Volume Backups are taken at a time chosen by the Data Center and may not meet the recovery requirements of some applications.

[Section 52](#) has details on restoring from Application Backups, the details will vary depending on options chosen by the application programmers.

If the same data sets have already been recovered by ABR full-volume recovery or ABR data set restore, the Application Restore will simply overlay those data sets with the contents from the Application Backup tape.

80.30 FDR/DSF/CPK/ABR I/O ERROR RECOVERY

TAPE I/O ERRORS FDR may encounter tape I/O errors. FDR always writes out highly blocked information to the tape. The records on tape are undefined format (RECFM=U) with lengths that vary from 20 bytes to 56K. Although tape cartridge drives are much less likely to encounter permanent errors than the "round reel" 3400-type drives, errors may still occur.

Most error recovery in cartridge drives is internal to the tape control unit, so when the error is reported back to MVS most errors are considered permanent. But normal access methods (usually BSAM) are used by FDR to read and write backup data sets, MVS's Error Recovery Procedures (ERPs) may be used to perform additional recovery. If this recovery is successful, FDR is not even aware that the error occurred. If the error is permanent, FDR may take additional actions which complement and extent the ERPs.

These notes mostly apply to backup data sets on disk as well, but disk I/O errors are extremely rare on modern disk subsystems.

TAPE ERRORS DURING DUMP During any kind of backup to a tape cartridge, FDR will issue message FDR200 to identify the error reported by BSAM, and will immediately terminate the backup (see the notes on SWAP below). Since the cartridge control unit and MVS ERPs cannot recover from the error, there is nothing more than FDR can do. The backup will need to be rerun with new output tapes. If the error repeats itself on the next run, you may need to ask your tape drive vendor to look at the error records and run diagnostics on the drives; since FDR is using IBM access methods (BSAM) there is little chance that FDR caused the error.

If you backup to round reels (3400-type drives), a permanent error will cause FDR to force a new output tape to be mounted and continue the backup on the new tape. Message FDR201 indicates that this took place. Up to 20 such errors are allowed before the backup is terminated (unless overridden by the MAXERR= operand). If the dump does complete, the backup can probably be restored successfully, but FDR, DSF, and CPK will terminate with a U0888 ABEND, and ABR will set a completion code 12 at termination, in order to draw attention to the errors (if this is not desirable, code TAPERRC=NO on the DUMP TYPE= statement).

TAPE ERRORS DURING RESTORE If a permanent I/O error is detected during a RESTORE or COMPAKT-from-backup operation, the program will issue message FDR200 or CPK502E to identify the error reported by BSAM. If the block in error is an FDR control record the program will immediately terminate. If the block in error is a data block, FDR will continue processing with the next block of data. A maximum of 20 blocks will be bypassed before terminating the restore unless MAXERR= is coded. The block in error may contain one or more tracks of data. These tracks will not be restored.

Even if a tape block is successfully read, its length is compared to an internal length field, in order to detect blocks which may have had undetected I/O errors, or which have been shortened by some other program. A block with a length error will be bypassed, just like a block with an I/O error. Tapes created by the FDR system should be copied only by the FDRTCOPY utility ([See Section 60](#)) or FATAR, a separate program product from Innovation. Because they contain blocks over 32K in length, utilities such as IEBGENER will not copy them correctly, resulting in block length errors.

An FDR366 or CPK582E message is issued at the end of the restore specifying the tracks, if any, that were lost. On a ABR data set restore, message FDR155 specifies the data set to which the missing tracks belong. FDR, DSF, and CPK will issue a U0888 ABEND and ABR will set a completion code 12 at the end of the restore to draw attention to the error, whether or not tracks needed for the restore were lost. If you code TAPERRC=NO on the RESTORE TYPE= statement, the error terminations will occur only if required tracks were really lost.

80.30 CONTINUED

TAPE SWAPPING MVS includes a facility, called SWAP, which attempts to recover from an error on a tape by swapping it to another drive and retrying the operation. It is possible to turn swapping off globally via the MVS console command SWAP OFF. Even if swapping is enabled, the operator will be asked for permission to do the swap and can designate the tape drive to which the tape will be swapped (IBM message IGF500D).

When a swap occurs, MVS must reposition the tape to the position it had when the error occurred and repeat the operation. For cartridge drives, which have buffers, all buffered WRITE data must be recovered and written to the new device.

On round reel (3400-type) drives, repositioning is done by counting the blocks that were already written or read. We have found this to be unreliable. For backups, we recommend that swapping not be used as it may result in lost data that is not detected until restore. For restores, at least one swap can be attempted; as noted above, FDR will inform you if any required data tracks were not found on the backup.

For cartridge drives, each installation should decide if it wants to allow swaps during backups. The repositioning on cartridges uses a hardware block ID which is much more reliable. Although the internal procedures required to swap a cartridge are more complex, it does seem to be reliable. But you must be aware that Innovation cannot guarantee that all the data we wrote was actually written to the tape. For restores, at least one swap can be attempted; more than one is probably futile.

Be aware that if a swap is apparently successful in recovering from the error, FDR is not informed that the error occurred and does not report it. You can't easily tell from a job listing that a swap occurred, because SWAP messages IGF500I, IGF502E, and IGF505I are not printed with the console messages in the job log at the beginning of the SYSOUT. Only the fact that the tape was mounted on one drive but dismounted from another will give you a clue about swapping.

DISK I/O ERRORS The FDR system uses its own CCW chains to read and write disk tracks. In many cases, MVS Error Recovery Procedures (ERPs) are allowed to recover from disk I/O errors. However, our own ERPs are often used in place of the system ERPs because of the unique nature of the I/Os we issue. Many errors are retried many times or in various ways in order to read or write the data if at all possible. If all recovery fails, the error will be reported in a message so that you can take appropriate action.

There are various I/O error messages, depending on where the error is detected and what type of I/O FDR is doing, but they are usually followed by diagnostic messages including a number of control blocks and other information, such as the IOB, DEB, DCB, UCB, and CCWs. The format of these can be found in various IBM manuals, but there are several significant pieces of information that may help you decipher the error condition:

- 1) the IOB contains 2 important fields: bytes 2 and 3 (last two bytes of the first word) contain sense data from the disk, which can be found in the appropriate control unit (such as 3880 or 3990) hardware manual; bytes 8-15 (the 3rd and 4th words) contain the CSW (channel status word), which is defined in the Principles of Operation manual for your system.
 - 2) in the CCW printout, the CCWs for the current I/O are printed in a vertical table or up to 4 on a line.
- Contact Innovation for assistance in diagnosing a disk I/O error.

80.30 CONTINUED

DISK ERRORS DURING DUMP Permanent I/O errors on disk while dumping usually result in disk tracks not being dumped, since the FDR system could not read the track having the error. Even if some data could be read from a track, a partial track will never be dumped. Most I/O errors will result in one disk track being bypassed, but some errors affect more than one track. When a track is bypassed because of an I/O error, a "dummy" entry is written to the backup tape in its place, so that a restore from that tape will know why it is missing (a warning message FDR150 is issued during a restore when such a missing track is encountered).

A maximum of 20 disk I/O errors can occur before the dump is terminated, unless the MAXERR= keyword is specified. If that maximum error count is not reached, FDR, DSF, and CPK will terminate with a U0888 abend, and ABR with a completion code of 12, to call attention to the errors; the backups will definitely be missing some tracks. The I/O error messages will contain a cylinder and track number, which you can use to determine which data sets are affected, by comparing it to a map of the volume.

I/O errors reading the VTOC or VVDS will usually make the backup unusable.

DISK ERRORS DURING RESTORE Permanent I/O errors writing to disk usually affect just one disk track. The track involved (identified by cylinder and track number in the I/O error message) may have partial data written to it, or none at all. Use that track ID to determine the data set affected.

INVALID TRACK FORMAT "Invalid Track Format" is an often-reported I/O error which is not really an I/O error. A disk track has a fixed maximum capacity, which varies by device type; the actual maximum data for a given track is determined by a formula based on the number and size of the records written to it. If an application or access method erroneously tries to write more data on a track than it will hold, the last record on the track is only partially present. During that write, the disk will indicate the "invalid track format" I/O error, but that partial last record may be left on the track. When that track is read by an FDR backup, the same "invalid track format" is reported by the disk. It is really reporting a logical data error, caused by a programming error, or sometimes an MVS bug.

"Invalid Track Format" can be recognized when the sense from the disk (the last 4 digits of the first word of the IOB in the FDR diagnostics) contain X '0040'. This error most frequently occurs in Partitioned Data Sets (PDSs). If so, it is possible that the error is not really in any real member, but is in dead space between members. The easiest way to tell is to copy the PDS with IEBCOPY; if no errors occur, the backup of the PDS is clean.

If you can be sure that the track getting the "invalid track format" does not contain any useful data, you can make the error go away with the IBM utility ICKDSF with the command:

```
INSPECT TRACKS(X 'cccc',X 'hhhh') NOPRESERVE
```

where cccc and hhhh are the cylinder and head of the bad track in hex.

80.31 ABR ABNORMAL TERMINATION PROCEDURES

This section discusses what should be done if an ABR Volume Backup, Archive Backup or restore of any kind abnormally terminates. The solutions are generalized and are not intended to cover every situation. Discretion should be used and if you are in doubt concerning what has occurred or the procedure to follow, contact Innovation for assistance.

First, be aware that there are two kinds of failures in ABR. ABR starts a subtask to actually perform the backup of a given disk volume or restore from a given backup file; for backups, if multiple TAPEX DDs are present, multiple subtasks are attached for different volumes. An ABEND in a subtask (indicated by message FDR319) will fail the backup or restore but ABR will attempt to continue with other backups or restores. But a lot of processing takes place in the ABR main task; an ABEND in the main task will terminate the entire step immediately, even though other backup subtasks may be running.

What should you do if an Archive Backup abnormally terminates? ABR does a fail-safe Archive. First it uses a subtask to backup all the selected data sets on one disk volume. After the backup has completed successfully, ABR will record these data sets in the Archive Control file. Finally, it scratches the data sets and either uncatalogs them or recatalogs for recall.

If the failure occurs while creating the backup, ABR will not record or scratch the data sets, but it will go on to process other volumes (it will call for a new scratch tape for the next backup, in case the error related to a tape error). You can recover by resubmitting the Archive Backup job or simply waiting until the next scheduled execution of Archive. Since data sets which were successfully archived were scratched, they will not be selected again.

If the failure occurs in the main task, the backup was complete and it is possible that some data sets were recorded and scratched before the failure. You must be sure that the backup tapes are retained, since they may contain the only copy of some data sets (since the tape was closed before the failure, your tape management system will probably retain it normally). Again you can either resubmit the Archive Backup or wait for the next regular Archive execution.

What should you do if there is an ABEND in the middle of an Volume backup? For full-volume backups, you should re-execute the backup for each volume that failed (subtask error indicated by FDR319 message) and every volume that was not processed because of a failure in the ABR main task. You may need to modify the backup job to include EXCLUDE statements for each successfully processed volume, e.g.,

```
EXCLUDE ALLDSN,VOL=vvvvvv
```

For incremental backups, ABR does not reset the update flag in the DSCB of a selected data set until the backup of the entire volume is complete and it is recorded in the ABR catalog, so the only exposure can occur if the failure occurs in the ABR main task while resetting those flags. If a failure occurs in a backup subtask, ABR will continue to process other volumes (it will call for a fresh scratch tape in case the failure was a tape error). You can simply resubmit the incremental Volume Backup job; it will only backup data sets from volumes that were not successfully processed in the previous backup.

What happens to pending requests if the system crashes while ABR is doing backups, Archives, or restores that were requested through the remote queue? Backups, Archives or restores that were completed before the crash are not affected. Backups, Archives or restores that were not completed will have to be re-requested by the users using the TSO/ISPF Panels or with program FDRABRUT.

Explanation: ABR reads the remote queue data set (ABRARCH for restores from Archive Backups, ABRREST for restores from Volume Backups, ABRARDQ for Archive Backups, ABRBKDQ for Volume Backups) at the beginning of the operation, constructs internal tables, and immediately resets the data set to empty. This minimizes the time that the data set is not available for users to add new requests. If the system crashes, ABR's internal tables are lost.

Note: The above answer always applies to restore requests, but it only applies to backup requests if the program FDRABRUT is processing them by adding control statements to the remote queue data sets. If FDRABRUT is processing backup requests by setting indicators in the DSCB, then if the system crashes while ABR is in the middle of a backup or ARCHIVE operation, the request will remain pending for the next backup run.

80.32 ABR AND TAPE MANAGEMENT SYSTEMS

**TAPE
MANAGEMENT
SYSTEMS**

ABR has no explicit interface to any tape management system, so ABR cannot directly control the retention and expiration of ABR-created tapes. However, ABR can indirectly influence the actions of tape management systems on ABR tapes by the expiration dates passed to the tape management systems when the tape files are originally opened by ABR, and by uncataloging tapes. This section describes the options that are available.

The tape management systems used most widely by ABR customers are

- CA---1 (also known as TMS and CA-ONE), from Computer Associates
- CA-DYNAM/TLMS (usually called just TLMS), also from Computer Associates
- DFSMSrmm (Removable Media Manager) from IBM.

ABR can be used successfully with these systems, with other tape management systems, or with no tape management software at all.

In this section, when we refer to TMS, we mean all tape management systems, not just CA-1.

There is information on Backup Retention and Tape Management in each of the 3 ABR sections of this manual

ABR Volume Backups ([Section 50.01](#))

ABR Archive Backups ([Section 51.01](#))

ABR Application Backups ([Section 52.01](#))

You should read these when you are designing the backup retention rules for a particular type of backup. This section provides additional information and relevant ABR options.

Note: FDR and FDRDSF, for manual full-volume and data set backups, have no interface and no influence on tape management. Backup tapes created by FDR or FDRDSF are managed just like any other tapes created by any application.

TMS OPTION

If you have any vendor's tape management system and also use ABR, you **must** let ABR know by activating the TMS option in the FDR global option table ([see Section 90](#), panel A.I.4.4). The TMS option has absolutely no effect unless you have ABR; ABR needs to know if there is tape management in the operating system.

With TMS enabled, ABR will change the techniques it uses to create LASTAPE files on tape (the LASTAPE option of ABR) and to call for a fresh scratch tape when the maximum file number on tape (MAXFILE=) is exceeded.

In the FDRARCH utility, TMS causes a REORG to uncatalog expired Archive Backups only if they are recorded as having an expiration date of 1999.000 (99000) which indicates catalog control in many tape management systems. However, if you want to uncatalog all expired Archive Backups, you can temporarily override the TMS option on a FDRARCH statement (DISABLE=TMS).

80.32 CONTINUED

**RETENTION
TECHNIQUES**

There are two techniques generally used by tape management systems for ABR tapes. The first is "date control", where an explicit expiration date is assigned to the tape data set. The other is "catalog control", where the data set is retained as long as it is still in a system catalog such as the ABR catalog.

Some other retention techniques supported by tape management may work with ABR, but another common tape management technique, "cycle control", will not work for ABR tapes. Cycle control only works for data sets which have the same name every time they are created; ABR backups are usually assigned unique data set names.

In all tape management systems, the assignment of an expiration date or retention technique to tape data sets can be controlled in several ways:

- 1) the common technique (supported since the early days of tape management) is to use the JCL keywords RETPD= (retention by number of days) and EXPDT= (for explicit expiration dates) to invoke date control. EXPDT= is also used to assign special retention techniques (e.g., EXPDT=99000 for catalog control) but this function may be disabled in some tape management systems.
- 2) some tape management systems now use other JCL keywords (such as ACCODE= for CA-1) to indicate retention techniques.
- 3) a "rules" data set can assign expiration information to a data set. This allows installations to implement tape management without having to update all of their JCL. The "rules" data set might even be able to override expirations set in JCL or by the program.

The details of these vary by tape management system. Please consult your vendor's documentation for details.

Methods (1) and (2) are usually used by ABR. They have the advantage that ABR knows the actual expiration date or retention method used to manage each backup file. In many cases, ABR saves and does processing based on that information.

Method (3) has the disadvantage that ABR does not know the actual expiration date or retention technique. If you use this method you may need to manually update ABR with the equivalent information, so it is not recommended.

**DATE
CONTROL**

When ABR creates a tape data set, an expiration date is passed to OPEN, where it is intercepted by your tape management system. Expiration dates may come from several sources:

- the user may specify a retention period via the RETPD= operand on an ABR DUMP statement. ABR adds that value to today's date to calculate the expiration date.
- The RETPD= (retention period) or EXPDT= (explicit expiration date) operands may be specified on the TAPEX and TAPEXX DD statements. For RETPD=, IBM adds the value to today's date to calculate the expiration date.
- If neither of the above, ABR will use a default RETPD= value to calculate the expiration date.

For all tape management systems, the tape will be retained until that expiration date is reached.

Date control can also be used on systems with no tape management software, since OPEN writes the specified expiration date in the header labels of each tape data. When no tape management system is involved, OPEN will not let a tape containing an unexpired data set as the first file be used for output unless the system operator replies to a console message.

Date control can be used with all ABR backup types, but remember that ABR will only record the date if it is specified on the ABR DUMP statement or the TAPE DD statements.

80.32 CONTINUED

PERMANENT RETENTION You may occasionally want to create an ABR backup which is to be kept permanently (at least until someone manually expires it). To do so, specify EXPDT=99365 or 99366 on the TAPEX or TAPEXX DD statement. All tape management systems will treat this as permanent retention, as will ABR itself.

Note: if a RETPD= value results in a calculated expiration of 99365, both MVS and ABR will automatically bump it to 2000.001 (January 1, 2000) to avoid confusion with permanent retention.

CATALOG CONTROL Catalog control will retain an ABR tape data set as long as the data set is still in a system catalog (such as the ABR catalog). [Sections 50, 51, and 52](#) explain the rules for cataloging each type of ABR backup.

For most tape management systems, catalog control is specified by EXPDT=99000 on the TAPEX or TAPEXX DD statement. Check your vendor's documentation for the setup necessary to make your tape management system accept EXPDT=99000.

When ABR has recorded the expiration of a backup as 99000, it will take special action for that backup. For example, when you do an FDRARCH REORG on the Archive Control File, it will uncatalog backups which no longer have any references in the ACF, but only if the recorded expiration is 99000. So, even if you are invoking catalog control by another means (such as ACCODE or a "rules" data set), you should also specify EXPDT=99000 so that ABR knows about it.

**UNCATALOG-
ING BY TAPE
MANAGE-
MENT
SYSTEMS** When you use date control with your tape management system, the TMS may uncatalog tape data sets when the tapes expire. This is usually an option, so you should consult the vendor's documentation. If at all possible, this option should be enabled so that obsolete ABR backups are uncataloged.

As explained in [Sections 50, 51, and 52](#), ABR also includes code to uncatalog backup tapes when they expire. Whichever system attempts first to uncatalog a given tape data set will succeed, and the other system will fail. Neither system will consider this to be a problem.

In some cases, ABR will not attempt to uncatalog a backup tape until long after it expired. For example, if you keep 5 generations of an ABR full-volume backup, but assign RETPD=14 to the daily incrementals, the incrementals will expire in 2 weeks but ABR will not uncatalog them until 3 more generations have been created (5 weeks total). These cataloged but expired incrementals will continue to appear on ABR reports, but any attempt to restore from them will fail. The solution is to enable the option allowing your TMS to uncatalog them.

80.32 CONTINUED

**MULTI-FILE
TAPES**

ABR usually puts more than one file onto each tape volume. In some cases, the tape files on one tape may have varying expiration dates.

The LASTAPE feature of ABR provides a method of improving tape utilization by allowing you to add tape files onto tapes that were created in a previous ABR step or job. LASTAPE is most useful if each ABR run uses a relatively small amount of tape, so it is most often used for Archiving. If you do not use LASTAPE, then each run of ABR will call for fresh scratch tapes. However, LASTAPE is likely to create tape files with varying expirations.

CA-1 will handle multi-expiration tapes automatically; it will not expire any tape volume until all files on that tape have reached their expirations (including catalog control).

CA-DYNAM/TLMS normally bases expiration of a tape volume on the CDS (Controlling Data Set), which is normally the first data set on the tape. For an ABR tape, this may have the lowest expiration on the tape, so the tape may expire long before data sets near the end of the tape reach their expiration. The ABR Installation Control Library (ICL) member TLMSXTRS contains instructions and code for a modification to the TLMS user exit TLMSXTRS, which will prevent ABR tapes from expiring until every file on the tape has expired. All CA-DYNAM/TLMS users running ABR should review this.

In CA-DYNAM/TLMS V5.5, there are new retention types (A, B, and C) which will cause tapes to be retained until all files on the tape have reached their expiration, similar to the way that CA-1 operates. These retention types are recommended for ABR-created tapes. When these retention types are used, the TLMSXTRS exit modification is no longer required.

WARNING: Both CA-1 and CA-DYNAM/TLMS require additional records in their master files to record tape data sets which are not the first data set on a tape (DSNBs in the TMC for CA-1, auxiliary records in the VMF for CA-DYNAM/TLMS). Since ABR is a major user of such data sets, some installations may need to increase the size of the master files or the number of such records when implementing ABR.

DFSMSrmm handles expiration of multi-file tapes individually by file, so it will retain each file until it reaches its expiration. A tape volume becomes scratch when all files on it have expired.

**MULTI-
VOLUME
MULTI-FILE
AGGREGATES**

ABR may also create multi-volume multi-file aggregates (sometimes called "tape sets") on tape. An ABR aggregate may have up to 4095 files and may occupy an indefinite number of tape volumes, although the number of tape volumes used for an individual backup file is limited to 19.

Both CA-1 and CA-DYNAM/TLMS handle such aggregates automatically. All of the volumes in an aggregate will be retained or expired as a unit, all at once.

DFSMSrmm, since it expires individual files, may return a tape to scratch status when all files on it have expired, even though it is part of a multi-volume aggregate. In OS/390 V2R10 and above, DFSMSrmm has an option to retain and expire multi-volume aggregates as a unit; you should enable this option for proper retention of ABR-created tapes.

80.32 CONTINUED

VARYING EXPIRATIONS Many installations wish to create monthly, quarterly, or yearly Volume Backups which are to be kept for a longer period than the normal ABR backup or to keep the full-volume backups for a longer period than their associated incremental backups. With a tape management system this is easy.

Each disk volume should be initialized to retain a number of ABR generations equivalent to the longest backup you plan to keep; this prevents ABR from uncataloging any backup which is still available for restores. For example, if you create weekly full-volume backups, and you plan to keep certain backups for one year, set each volume to keep 52 or 53 generations.

Specify the proper retention on each ABR backup job. To expire daily incrementals early, specify the proper RETPD=nn in the daily job, either on the DUMP statement or the TAPE DD statement (we recommend a minimum retention of 14 days or two generations). Your normal full-volume backup job can specify the usual retention (e.g., 35 days for monthly retention), but special full-volume backup jobs can replace the normal job when you want to create long-term retention backups, specifying the proper RETPD= value. Alternately, you can use a tape management utility to extend the retention of certain backups. ABR will not care when they expire as long as they are still cataloged; the tape management system will uncatalog backups which reach their expiration.

ABR can restore data sets from any backup which is still in the ABR catalog. However, the OLDBACKUP option of ABR will only track the most recently created backups of a data set (up to 14), even if they have expired, so recovery from older backup tapes may require reviewing the ABR dump listings, or FDRABRM PRINT BACKUP reports ([See Section 53](#)) to find the required backup of a specific data set, and specifying the VOL=, GEN=, and CYCLE= options on the SELECT statement or on the ABR ISPF panels.

VAULTING How do you tell the Tape Management System to select ABR backup tapes for offsite vaulting?

ABR backups always have a specific naming convention. It is described in detail in [Sections 50](#), [51](#), and [52](#), but briefly it is:

FDRABR.Vvvvvvv.Cnggggcc	for Volume Backups
FDRABR.Vvvvvvv.Bnyyyyyyx	for Archive Backups
userindx.Vvvvvvv.Bnyyyyyyx	for Application Backups

The FDRABR high-level index can be changed during the installation of ABR. Application Backups will use an application-specified index in place of FDRABR. "vvvvvv" is the volume serial of the input disk, and "n" is the copy number (1 through 9).

In general, you will want to send COPY2 of most backups to the offsite vault. The details of how to do this vary by tape management system; consult the vendor's documentation.

Most tape management systems will allow you to specify a data set name prefix for vaulting selection. For ABR backups you may need to define a rule for backups from each disk volume, e.g.,

FDRABR.VPROD01.C2

FDRABR.VPROD02.C2

Change C2 to B2 to select Archive Backups.

If your TMS supports masking for vault selection, you may be able to specify something like:

FDRABR.V!!!!!!C2

to select COPY2 of any Volume Backup (full and incremental) for any volume.

CA-1 V5.1 and above supports masking. For CA-1 V4.8 through V5.0 only, the ABR Installation Control Library (ICL), member TMSMASK contains a user-submitted source modification to the CA-1 vaulting program which allows it to accept a "masking" character (!) which will always match, so that "V!!!!!!" will match on *all* disk volumes, and only one DSN= statement is required.

80.32 CONTINUED

**ONLY FULL-
VOLUME
BACKUPS
OFFSITE**

Many installations decide not to devote the resources that it takes to keep offsite copies of all ABR backups. They may decide that if a disaster happens, it will be acceptable to recover to the status as of the beginning of the current week. A simple way to set this up is to make duplicate copies of only the full-volume backups, and not the incremental backups, so the vaulting controls shown above will select only the full-volume backups. Since some tape management systems allow selection based on creating jobname (jobname qualification), another option is to run the full-volume backups under a different jobname than the incrementals.

COPY1=COPY2

An installation may want to create just one copy of all of the backups, both incremental and full-volume, but send offsite only the full-volume backups. To do this, they can specify the operand COPY1=COPY2 on the DUMP Command for only the full-volume backups:

```
DUMP    TYPE=FDR , COPY 1=COPY 2    . . .
```

COPY1=COPY2 tells ABR that even though the run is creating only one copy, the backup data set names should be in the form FDRABR.Vvvvvvv.C2ggggcc instead of FDRABR.Vvvvvvv.C1ggggcc. The vaulting system would still vault all COPY2 tapes created by ABR backups, but these tapes would only include the full-volume backups (and there would not be any COPY1 for these backups).

80.33 STORAGETEK PERFORMANCE PRODUCTS

StorageTek offers two products which can be used to enhance the operation of FDR:

- HSDM (High Speed Data Mover) – a hardware option which allows FDR to backup internal compressed track images
- ExHPDM (High Performance Data Mover) – a software product which can consolidate multiple tape outputs (such as FDR backups) onto one tape, improving the utilization of high-performance tapes.

Note that HSDM and ExHPDM can be used together in FDR, providing the benefits of both products.

More information on HSDM and ExHPDM can be obtained from StorageTek, or visit their web site at:

www.storagetek.com

HSDM HSDM (High Speed Data Mover) is very simple: it allows FDR to read and write disk tracks in a compressed format, significantly reducing the elapsed time for backups and restores. It is an extra-cost hardware option available on the StorageTek SVA (Shared Virtual Array) disk subsystem and on some IBM RVAs (RAMAC Virtual Array).

The SVA (and the RVA) store data tracks internally in a compressed format. Even track images in cache remain in a compressed format until they need to be sent to the host.

HSDM allows the host to read and write the compressed image of a data track, using special channel programming. This avoids compression overhead and reduces the amount of data that must be sent over the channel. It also reduces the amount of data to be written to the backup data set.

Support for HSDM has been implemented in FDR, FDRDSF and FDRABR, as well as SAR restore, but only if you are licensed for FDR InstantBackup. In any FDR backup step, add the operand DCT=YES (DUMPCOMPRESSEDTRACK=YES) to the DUMP statement. This will invoke HSDM for any input disk that has the option installed. Disks that do not have the option installed or are not capable of HSDM (not a SVA or RVA) will be dumped normally.

During a backup of an HSDM-capable disk, with DCT=YES specified, FDR will read the compressed track images from the disk subsystem and write them to the backup. Tests have shown up to a 60% reduction in elapsed backup time when HSDM is used.

Most tape cartridge drives support hardware compression (called IDRC on some drives) and it is usually the system default to use it. Since HSDM data is already highly compressed, the tape hardware compression is unlikely to significantly compress it further. The overhead of the hardware compression may actually slow down the backup. So, it is recommended that you specify TRTCH=NOCOMP on the TAPEX DD statements used with HSDM backups to suppress hardware compression. However, if you will be backing up a combination of HSDM-capable and non-HSDM disks in the same step, you may need to leave compression enabled to compress the non-HSDM backups.

During a restore, FDR will automatically recognize that a backup tape contains compressed HSDM track images. If the output disk is also HSDM-capable, the compressed track images are restored directly, improving restore elapsed time. If the output disk is not HSDM-capable, FDR will use a software decompression routine to reconstruct each original uncompressed track image and write it back to the disk with standard write CCWs.

Note: the FDR COMPRESS= option will be ignored if DCT=YES is specified, because the data is already compressed.

80.33 CONTINUED

ExHPDM ExHPDM (High Performance Data Mover) is a StorageTek software product which consolidates multiple output data streams (such as concurrent FDR backups) into one tape file, interleaving the data from the data streams. It exploits modern high-performance, high-capacity tape drives such as the StorageTek Redwood, Timberline and 9840. However, it is hardware-independent and also supports the IBM Magstar 3590.

It gets its performance improvements in two ways:

- it writes large blocks of data, up to 256K in length
- by combining multiple output streams into one tape file at a higher total data rate, it improves the efficiency of the tape.

ExHPDM operates as a permanent MVS started task (STC) and is also defined as a MVS subsystem. Any program which uses BSAM or QSAM to write/read a sequential file can use ExHPDM, although it is most useful with applications which write large quantities of data from several sources, such as concurrent FDR backups.

To invoke ExHPDM, the user must add the SUBSYS= operand to the DD statement for the sequential data set, for example,

```
//SYSUT2 DD DSN=XYZ,DISP=(,KEEP),SUBSYS=(SOV,'additional parameters')
```

SOV is the name of the ExHPDM subsystem; this is the default name provided with ExHPDM, but your installation may have changed it. You might even have several ExHPDM subsystems operating. Specify the proper subsystem name in the SUBSYS= operand.

For output data sets, you may need to specify additional parameters which instruct ExHPDM how to manage this data set. Some of the parameters, such as CLASS, relate to other definitions in the ExHPDM startup parameters which allow it to determine the type of tape to use and how many such output streams can be combined into one tape. You can also code selection rules in ExHPDM for the automatic assignment of classes to output data sets.

For input data sets, only SUBSYS=SOV is usually required.

ExHPDM will allocate the required input and output tapes in its own address space. The output of the user job will not see any tape mounts, but it will receive ExHPDM messages in SYSMSG (system messages) indicating that ExHPDM was used.

The data set names specified by the user are called "client data sets" by ExHPDM; these names are recorded in the ExHPDM database and are usually cataloged to a dummy volume serial of "***SOV*". The interleaved tape data sets created by ExHPDM are called "stream files"; their names are specified in ExHPDM startup parameters. Stream files cannot be read directly by client programs.

The ExHPDM user documentation explains all of the above in much greater detail.

80.33 CONTINUED

USING ExHPDM WITH FDR ExHPDM is supported for programs FDR and FDRDSF as of V5.3 level 30; it is also supported for FDRTCOPY when used to copy FDR and DSF backups. In FDR and FDRDSF, you can backup several disk volumes concurrently and invoke ExHPDM to intermingle them in a single tape file. The concurrent backups can be in separate FDR jobs, or in one job using the ATTACH or MAXTASKS= options.

When used with FDR and FDRDSF, the level of ExHPDM must be at least Version 1.1 plus PTF L1P013W.

To direct an FDR or FDRDSF backup to ExHPDM, use JCL similar to:

```
//TAPE1 DD DSN=BACKUP.VPROD01(+1),UNIT=TAPE,
//      SUBSYS=(SOV,'CLASS(FDRBKUP)'),DISP=(,CATLG)
```

The optional additional parameters, such as CLASS, pass additional information to ExHPDM about the management of the backup. In this example, FDRBKUP is a class name your installation has defined to ExHPDM which controls the type of tape to be used for output and how many concurrent backups are to be put to the same tape, as well as other ExHPDM characteristics. More information and examples are found in the ExHPDM user documentation.

ExHPDM will internally catalog the backup data set to a dummy volume serial of "***SOV*" and a device type of X'00000000". ExHPDM maintains its own data base with records of the location of every data set under its management. The ExHPDM SOVADMN utility may be used to obtain information about the stream file containing a given client data set and the real tape volumes on which it resides.

To restore a backup that was sent to ExHPDM, use JCL similar to:

```
//TAPE1 DD DSN=BACKUP.VPROD01(0),DISP=OLD,
//      SUBSYS=SOV
```

The SUBSYS= operand is required to read a data set which was written to ExHPDM, which will locate and mount the actual tape required. ExHPDM will find your data set name in its records, mount the proper tape, and extract the data records that belong to your data set from the tape.

To get the benefit of ExHPDM, you must do concurrent backups. These backups can be in one FDR step, using the ATTACH or MAXTASKS= operands with multiple TAPEX DD statements, or you can run separate FDR backup jobs concurrently. In either case, ExHPDM will combine the backups into one or more stream files (depending on your ExHPDM parameters).

Tests at Innovation indicate that if you add ExHPDM to backups which are already running concurrently to multiple output tapes, ExHPDM will reduce the number of tape drives and tape volumes required. If you are running backups serially to files stacked on tape or to separate tape volumes, ExHPDM will save a significant amount of elapsed time.

You can restore from any of the FDR backups that were intermingled on the ExHPDM tape at any time; ExHPDM will extract the proper data. This may be somewhat slower than a restore from a standard FDR backup because it has to skip over data for the other backups. You can also do concurrent restores from backups that were intermingled and ExHPDM will present each restore with the appropriate data blocks.

You will find ExHPDM JCL examples in [Sections 10](#) (FDR), [20](#) (DSF), and [60](#) (FDRTCOPY); there are also examples in the ExHPDM user documentation.

80.33 CONTINUED

**USING
ExHPDM WITH
ABR**

ExHPDM is supported with ABR volume backups and application backups (FDRAPPL) as of V5.3 level 40. In addition, FDRTSEL and FDRTCOPY have been enhanced to use ExHPDM. In ABR, you can backup several disk volumes concurrently and invoke ExHPDM to intermingle them in a single tape file.

Note: you must be licensed for FDR InstantBackup to be able to use ExHPDM with ABR. The level of ExHPDM must be at least Version 1.1 plus PTF L1P015A.

ExHPDM is not supported for ABR Archive backups. Data sets are usually recalled from Archive backups one at a time, so ExHPDM would tend to elongate restore times.

ABR automatically dumps volume concurrently, based on the number of TAPEX DD statements in the ABR step and the number of disk volumes selected for backup. To direct an ABR backup to ExHPDM, use JCL similar to:

```
//TAPE1 DD SUBSYS=(SOV,'CLASS(ABRBKUP1)')
```

The optional additional parameters, such as CLASS, pass additional information to ExHPDM about the management of the backup. In this example, ABRBKUP1 is a class name your installation has defined to ExHPDM for COPY1 backups which controls the type of tape to be used for output and how many concurrent backups are to be put to the same tape, as well as other ExHPDM characteristics. More information and examples are found in the ExHPDM user documentation.

If the backup must be cataloged, ABR will internally catalog the backup data set to a dummy volume serial of "***SOV*" and a device type of X'00000000". For FDRAPPL application backups, that volser and type are also recorded in the Application Control File. ExHPDM maintains its own data base with records of the location of every data set under its management. All ABR reports will show the **SOV* volser. The ExHPDM SOVADMIN utility may be used to obtain information about the stream file containing a given client data set and the real tape volumes on which it resides.

If you use the DYN TAPE option during an ABR restore, ABR will recognize the special volser and will allocate the backup file with SUBSYS=SOV. This is the default subsystem name used by ExHPDM; if your installation uses a different subsystem name, please contact Innovation for assistance. ExHPDM will then mount the tapes required for the restore.

80.33 CONTINUED

CONSIDERATIONS FOR EXHPDM USE

Here are rules and guidelines for the use of ExHPDM with FDR:

- 1) Do specify MAXERR=1 when the input or output is directed to ExHPDM, causing the operation to fail on the first "tape" error. Unlike real tapes, an error from ExHPDM is not recoverable so there is no point in attempting to continue.
- 2) ExHPDM records the expiration date passed when a client data set is opened. This expiration may be specified by the EXPDT= or RETPD= operands on the output DD statement. In ABR, it can also be specified by the RETPD= operand on the DUMP statement or ABR defaults. However, these expirations are not seen by your tape management system. ExHPDM will set the expiration date of the stream file (which may contain client data sets with varying expirations) according to ExHPDM initialization parameters. Innovation recommends that you setup ExHPDM to expire a ABR stream file when all client data sets in it have expired (see the ExHPDM documentation for details) since this makes the expiration passed by ABR meaningful.
- 3) The ABR LASTAPE option will not work with ExHPDM. ExHPDM does not allow data to be added to a stream file once the tape has been dismounted. However, the normal operation of ABR will dump many disk volumes in a single step; ExHPDM will combine them into one or more stream files. You may need to add the RETAIN operand to the STREAM definition in ExHPDM startup parameters so that an output tape is retained for a period of time to allow additional client data to be added to it.
- 4) Concurrent data set restores (FDRDSF or FDRABR) can be executed only by running multiple concurrent restore jobs. However, you must not do so if pieces of multi-volume data sets might be restored from 2 or more backups concurrently; FDR requires that pieces of a multi-volume data set be restored one at a time.
- 5) FDR cannot read an ExHPDM tape directly. To restore with ExHPDM, the ExHPDM subsystem must be active. Likewise, SAR can never read an ExHPDM tape. If you are creating backups for use at a disaster site or a new data center, any disk volumes that you intend to restore with SAR and those that may be restored before ExHPDM can be activated, must be dumped to normal FDR backup tapes. See the ExHPDM user documentation for disaster/recovery considerations.
- 6) In a FDR or FDRDSF backup step, the number of TAPEX DD statements exactly specifies the number of disk volumes that will be backed up in the step, but in an ABR backup, the TAPEX DDs specify the number of concurrent backups, but the number of disk volumes that will be backed up may be much larger. You should consider the effects of putting too many backups on a single stream file since it may elongate restore times.
- 7) Contact Innovation if you want to use FDRCLONE or FDRDRP ([Section 50.70](#)) with ExHPDM backups.

RESTORE FROM EXHPDM

All FDR, DSF and ABR restores are serial restores, meaning that no matter how many restores you request in a single FDR restore job, only one backup data set will be read at a time. For full-volume restores this means that only one disk volume will be restored at a time. For data set restores, only the data sets that were backed up from a single disk volume may be restored concurrently.

Since ExHPDM intermingles data from multiple disk backups in a single stream file, normal FDR restores from ExHPDM tapes will not be efficient since each restore task must skip over data belonging to other intermingled disks. So restores from ExHPDM may take longer than equivalent restores from normal tape unless additional steps are taken.

You can execute concurrent FDR and ABR full-volume restores by executing multiple such restore jobs in parallel, each job restoring just one disk volume. If the backups that are intermingled on the same stream file are restored concurrently, ExHPDM will be able to restore them efficiently, making one pass of the data to restore all of the disks involved.

Note: DSF and ABR do not support concurrent data set restores if multi-volume data sets are involved. Do not attempt to improve data set restore performance by submitting restores from multiple disks in parallel unless you can be sure that no multi-volume disk data sets are to be restored.

Experience has shown that the easiest way to achieve efficient ExHPDM restores is to submit separate restore jobs for each of the disk volumes involved and let them execute concurrently (be sure that there are enough MVS initiators to execute these jobs concurrently). ExHPDM will mount as many stream files as it can and will execute the restores for the disks whose data is intermingled on those stream files. Restores whose data resides on tapes which cannot be immediately mounted will simply wait until tape drives are available to mount them.

80.90 FDRCHSM CONVERSION AID

Included with the FDRABR product is a utility to automatically convert data sets migrated under DFHSM to FDRABR archived data sets. Installations seeking the outstanding performance, reliability, and features of the FDRABR DASD management system with data now under DFHSM control will be able to safely move the data to FDRABR control. The conversion aid runs as a background job, and is designed for incremental conversion to ABR, selecting a user controlled number of migration level two (ML2) tapes for each increment. This allows for the conversion to be done as time and resources permit. User control of the process is provided for by a simple control card input, and optional system operator interaction. All DFHSM interactions are done using documented interfaces, with no dependencies on DFHSM internals. The conversion aid does do extensive validation to insure that data integrity is maintained, and a restart capability is provided in case of abends or system failures while the conversion aid is running.

The conversion is done by recalling all of the migrated data sets on a selected ML2 volume, then invoking FDRABR to archive those data sets. The process can continue automatically for a user specified number of ML2 volumes during one execution of the conversion aid. Also, if resources permit, multiple conversion runs can be done concurrently. The conversion aid has no dependencies on DFHSM internals, as everything is driven based on the documented DFHSM external commands and reports.

The conversion process is intended to be done over a planned period of time. The amount of time required for the conversion is dependent on the amount of data to be converted to FDRABR control, and the available resources within the environment. The conversion process should be started slowly, only doing a few ML2 tapes at a time, so that experience in the operation of the utility can be gained, and estimates of how long to do each tape in the environment can be determined. After some experience is gained, the conversion can proceed at the desired pace.

All of the interaction with DFHSM is through the use of the DFHSM TSO command interface. The conversion aid builds sets of TSO commands to be issued in a sequential data set, then automatically invokes TSO in batch mode to execute the commands. This process completely eliminates any dependencies on the internals of DFHSM, its control files, and the format of the data that is to be converted. The conversion aid is dependent on the externals, the documented commands and certain report formats.

After validating the environment and the JCL, the conversion aid issues a request to DFHSM for an inventory of all of the ML2 tapes. Using that list, and the specified criteria on the control statement input, the conversion aid selects an ML2 volume to be converted. Once an ML2 volume is selected for processing, the conversion aid then requests an inventory of the migrated data sets on the selected ML2 volume. From the inventory of data sets on the ML2 volume, DFHSM TSO RECALL commands are built, directing the recalls to the DASD volume specified by the user, and invokes the batch TSO processor to issue the RECALL commands.

Upon completion of all of the RECALL commands, the conversion aid validates that each data set was properly recalled. For all data sets that were properly recalled, control cards are built for those data sets to be archived for automatic recall by FDRABR, and FDRABR is invoked to archive the data sets. Once the FDRABR archival process is complete, then if so specified by the user, another ML2 volume is selected and processed in the same manner.

The conversion process is a very safe way to help customers move data sets migrated by DFHSM to an archived FDRABR format. The design of the program is completely oriented towards data integrity, and as such is not a "fast" process. The conversion aid does have a restart capability, should there have been a failure during the recall or archive process, to resume processing. The conversion aid also has a simulation mode. Use of this mode will validate that the conversion aid is working with DFHSM, and provides the user with a list of ML2 volumes that will be processed.

Detailed documentation for this utility is in member FDRCHSM in the FDR ICL (Installation Control Library).

THIS PAGE WAS INTENTIONALLY LEFT BLANK

90.01 FDR INSTALLATION INTRODUCTION

FDR should be installed using the interactive installation procedure documented here in Section 90. This requires TSO with ISPF.

Innovation strongly recommends the use of this interactive installation procedure. It is easier to use and less prone to errors than any manual procedures. However, if TSO or ISPF is not available or other considerations prevent use of the interactive procedure, please follow the alternate installation procedure documented in [Section 91](#).

**DISTRIBUTION
TAPE FORMAT**

FDR is distributed on a standard label (SL) magnetic tape volume. The volume serial and status of the distribution tape will be clearly marked on the external tape label. The volume serial of the tape will be:

FDR53P - a production (licensed) version of one or more components of the FDR product family. The particular components on your tape are indicated on the external tape label and depend on your license from Innovation Data Processing.

FDR53T - a trial (evaluation) version of FDR (always contains all FDR components). The trial programs will expire (stop functioning) on the expiration date indicated on the external label.

The files on the tape are:

- File 1: DSN=SAR
an IPLable version of SAR (Stand-Alone Restore), the stand-alone component of FDR.
- File 2: DSN=LOAD
IEBCOPY-unloaded library of programs.
- File 3: Empty. Formerly contained documentation.
- File 4: DSN=ICL
IEBCOPY-unloaded library of installation jobstreams and supplementary documentation.
- File 5: DSN=CLIST
IEBUPDTE-format file of ISPF dialog CLISTs.
- File 6: DSN=PANELS
IEBCOPY-unloaded library of ISPF dialog Panels.
- File 7: DSN=MESSAGES
IEBCOPY-unloaded library of ISPF dialog Messages.
- File 8: DSN=SKELETON
IEBCOPY-unloaded library of ISPF dialog Skeletons.
- File 9: DSN=FDR.INSTALL
an executable program which is used to load the files on this tape to disk.
- File 10: DSN=JCL
IEBCOPY-unloaded library of all example JCL from this user manual.

An IPLable copy of STAND-ALONE RESTORE (SAR) is in the first file of the tape. Since the distribution tape is labeled (SL), the IPL function must be performed 5 times to load SAR from this tape (the first 4 attempts will get IPL errors as the system reads the tape header labels; the 5th attempt will be successful). For a more convenient IPL, you can create an unlabeled tape containing SAR, or put a copy of SAR on one or more disk volumes. [See Section 15.20](#) for details.

A CD-ROM is also included with the FDR product. The CD-ROM contains a copy of the entire FDR user manual in IBM BookManager format and also in Adobe PDF format for ease of printing.

90.01 CONTINUED

NEW INSTALLATIONS

Use the following checklist as you proceed through the installation. The PRODUCT column indicates to which FDR components the step applies; you need be concerned only about the steps for the components for which you are licensed. Trial customers will have all components of the FDR system; we hope that you will test all the pieces of FDR, but if necessary you can skip the installation steps for components you will not be using.

CHECK	STEP	PRODUCT	SECTION	DESCRIPTION
	1	ALL	90.03	Load all files from the FDR distribution tape
	2	ALL		Review the NEWS member in the Installation Control Library (ICL) for the latest news and information
	3	ALL	90.10	Invoke the FDR Install ISPF Dialog
	4	ALL	90.11-13	Set the FDR Global Options
	5	CPK	90.14	Set the COMPAKTOR Global Options
	6	ABR	90.15-19	Set the ABR Global Options
	7	ABR, FDREPORT	90.20-21	Set reporting Global Options
	8	ABR, FDRREORG	90.22	Set the Operating System exit Global Options
	9	FDRREORG	90.23	Set the FDRREORG Global Options
	10	CPK	90.30	Set up the COMPAKTOR Unmovable Table
	11	FDRREORG	90.31	Set up the FDRREORG NOREORG List
	12	ABR	90.32	Set up the ABR Protect and Restore Allocation Lists
	13	ALL	90.40	Authorize the FDR Programs
	14	ABR	90.41	Create the ABR Catalog
	15	ABR	90.42	Update DASD VTOCs with ABR Disk Processing Options
	16	ABR	90.43	Create the ABR Archive Control File
	17	ALL	90.44	Set FDR Dialog Options
	18	ALL	90.45	Create the ABR Remote Queue Data sets
	19	ALL	90.46	Install the FDR ISPF dialogs in production
	20	FDR	90.47	Load SAR onto disk
	21	ALL	90.50	Test online and batch functions
	22	ABR, FDRREORG	90.60-62	Install and test MVS exits

90.01 CONTINUED

EXISTING INSTALLATIONS

Use the following checklist as you proceed through the installation. The PRODUCT column indicates to which FDR components the step applies; you need be concerned only about the steps for the components for which you are licensed. Although you will be copying most of the options and customization from your previous FDR release, you may need to review other installation steps (as shown in the "New Installations" checklist above) to be sure that all options are set appropriately. Also, existing customers who are installing new FDR components may need to review the "New Installations" checklist to insure that all steps relating to the new components are completed.

CHECK	STEP	PRODUCT	SECTION	DESCRIPTION
	1	ALL	90.02 or 90.03	Load all files from the FDR distribution tape
	2	ALL		Review the Summary of Modifications (section 01.04), and the NEWS member in the Installation Control Library (ICL), for the latest news and information
	3	ALL	90.10	Invoke the FDR Install ISPF Dialog
	4	ALL	90.11	Copy Global Options, COMPAKTOR Unmovable Table, ABR protect and restore allocation lists, and FDRREORG NOREORG list from previous version
	5	ALL	90.44	Copy FDR Dialog Options
	6	FDR	90.47	Load SAR onto disk
	7	ALL	90.50	Test online and batch functions
	8	ABR, FDRREORG	90.60-62	Install and test MVS exits
	9	ALL	90.46	Install the FDR ISPF dialogs in production

**TRIAL
CUSTOMERS
INSTALLING
PRODUCTION**

If you have installed and tested a trial version of FDR, and are now installing the production installation tape you obtained when you signed the contract with Innovation, you can follow the above checklist ("Existing Installations"). This allows you to copy all of the options and customizations you set during testing, but you may want to review your options (especially ABR data set names) to be sure they are appropriate for your production use of the products.

90.01 CONTINUED

FDRCLONE [Section 90.24](#) describes options that affect FDRCLONE. These options will not be set during the normal installation of FDR. However, when you use FDRCLONE to "clone" data sets on a test system or at a disaster/recovery site, you must enable the FDRCLONE option on only on that system.

**VIEWING/
PRINTING THE
SOFTCOPY
DOCUMENTATION**

A copy of this FDR manual is provided on a CD-ROM or diskette in softcopy format for use with IBM BOOKMANAGER products, including

- Bookmanager READ/MVS (a standard component of OS/390)
- Bookmanager READ/DOS
- Bookmanager READ/2 (for OS/2)
- Bookmanager READ for Windows
- Bookmanager READ/6000 for AIX/6000

The CD-ROM also contains a copy of the IBM Library Reader, which enables you to use the Bookmanager documents without being licensed for one of the IBM products.

In addition, the CD-ROM contains the manual in Adobe Acrobat PDF format, plus a copy of the Adobe Acrobat Reader, enabling you to view and print the manual in a format identical to the printed manual.

The manuals on the CD-ROM can be used directly on any of the supported systems except MVS. To use Bookmanager files on MVS, you must have READ/MVS and you must upload the Bookmanager files to MVS disk storage, following directions on the CD-ROM..

90.02 RE-INSTALLATION OF THE FDR DISTRIBUTION TAPE

EXISTING CUSTOMERS: If you have V5.3 of the FDR ISPF dialogs already installed from an earlier level, or from a trial tape, you have the option of loading the FDR distribution tape with the currently installed FDR Install dialog, ISPF option A.I.1 (unless instructions included with your tape indicate that this is not possible due to tape format changes). Or you may use the interactive install program documented in [Section 90.03](#).

From the FDR Installation Options Menu (usually accessed by A.I), enter 1 (INSTALL) to display this panel:

**PANEL A.I.1
RE-INSTALL
FDR TAPE**

```

----- FDR RE-INSTALLATION -----
COMMAND ==>

      SUBMIT - SUBMIT INSTALLATION JOB      EDIT - EDIT INSTALLATION JOB

      TAPE
      SEL DATA SET  DESTINATION DATA SET NAME (QUALIFIED)  VOLUME SERIAL  NEW/  BLOCK
      -----
      S  CONTROL    IDP.ICLFDR53                                NEW  3120
      S  LOADLIB    IDP.MODFDR53                                NEW  6144
      S  CLIST      IDP.DIALOG.CLIST                             NEW  3120
      S  PANELS     IDP.DIALOG.PANELS                           NEW  3120
      S  MESSAGES   IDP.DIALOG.MESSAGES                         NEW  3120
      S  SKELETON   IDP.DIALOG.SKELETON                         NEW  3120
      S  JCL        IDP.JCLFDR53                                NEW  3120

TAPE VOLUME SERIAL ==> FDR53P      TAPE UNIT NAME      ==> TAPE

JOB STATEMENT INFORMATION:      SYSOUT CLASS      ==> *
==> //useridA JOB (ACCOUNT),'NAME',
==> //          NOTIFY=userid
==> // *
==> // *

```

Warning: if you are installing a trial tape (serial FDR53T) but you are already licensed for some components of FDR, be sure to install the tape into a set of libraries totally different from those used for the production components. All components on the trial tape, including those you are already licensed for, will be trial (date-protected) versions.

All the data sets which can be loaded from the distribution tape are displayed. The SEL column has an "S" preset for all of them; you can blank out the S for any data set that you do not want to load.

For those that you do select, you can modify the data set name column to any name you desire. When you press ENTER, the OLD/NEW column will indicate if the name you selected already exists (OLD) or will be allocated (NEW).

If OLD, members will be added to the specified library.

If NEW, the blocksize that will be used for the allocation is shown. You should specify a volume in the VOLUME SERIAL column, unless the volume will be chosen by SMS; otherwise, the data sets may be allocated on any STORAGE volume in your system.

Note: during the installation, the ISPF dialogs will be modified to use the data set names you specify here. If you later rename the FDR libraries or copy the members to other libraries, you can update the dialogs with ISPF panel A.I.1A ([see Section 90.44](#))

**TAPE
SPECIFICA-
TIONS**

Enter the volume serial of the distribution tape (FDR53P for a production tape or FDR53T for a trial tape) and also a unit name which will allocate a tape drive capable of reading the distribution tape (a value which is valid for UNIT= in JCL).

90.02 CONTINUED

**JCL
INFORMATION**

Since a batch jobstream will be generated to load the tape, specify SYSOUT CLASS and JOB statement information appropriate for running the job.

Once the specifications are complete, you can enter, on the COMMAND line, SUBMIT to submit the job immediately to load the tape, or EDIT to edit and change the job before submission (or to save it).

While loading the FDR CLIST library, the jobstream will update all references to the various FDR libraries to the names you have specified here. When you use the newly loaded FDR ISPF dialogs, all the proper library names will be preset for you. If you later change the names of the libraries, you must update the dialogs as shown in [section 90.46](#) in the topic "Changing the Names of the FDR Libraries".

A list of the load modules used by each component of FDR can be found in the ICL (Installation Control Library) in member MODULES.

90.03 LOADING THE FDR LIBRARIES FROM THE DISTRIBUTION TAPE

The FDR Tape Install program, FDRLOAD, makes the installation of an FDR distribution tape very easy. You can execute FDRLOAD directly from tape if you have access to a TSO userid that has the "MOUNT" attribute, or if you are able to issue or request a command on a system console to have a tape mounted. Otherwise, you must copy the Tape Install program to disk using the JCL shown in Step 3.

If you have access to a TSO userid with the MOUNT attribute, logon to that id and proceed to Step 2. If you don't know if your userid has the MOUNT attribute, you probably don't so proceed to step 1.

STEP 1 Use this step if your TSO userid does not have MOUNT privileges and you are able to issue (or request to be issued) a command on a MVS system console to have a tape mounted. If your userid has MOUNT privileges, proceed to Step 2; otherwise proceed to Step 3.

If your TSO userid does not have the MOUNT attribute, you can still access a tape from TSO by having the operator issue a MOUNT command. You or the operator must mount and ready the tape on a free tape drive **BEFORE** issuing the following command on an MVS system console:

MOUNT *uuu*,VOL=(SL,FDR53T)

Change "*uuu*" to the actual tape unit address.

Change **FDR53T** to **FDR53P** if this is a production tape.

Now go to Step 2, but remember that when you are done with the tape, it must be unloaded by the MVS system command:

UNLOAD *uuu*

MOUNT /1234,VOL=(SL,FDR53T) and UNLOAD /1234

Note: if the tape unit has a 4-digit address, you must precede the address with a slash on the MOUNT command, and may also do so on the UNLOAD command, e.g.,

STEP 2 Use this step if your TSO userid has the MOUNT attribute, or if you have completed Step 1.

If you are using ISPF, issue the following TSO commands from ISPF Option 6 (TSO COMMANDS). You can also exit ISPF and issue them from the TSO "READY" prompt.

Enter this TSO command to allocate the FDR distribution tape:

ALLOC DA('FDR.INSTALL') VOL(FDR53T) UNIT(*tape*) POS(9) SHR

Change "*tape*" to an appropriate tape unit name..

Change **FDR53T** to **FDR53P** if this is a production tape.

If you get the message "IKJ56221I DATA SET FDR.INSTALL NOT ALLOCATED, VOLUME NOT AVAILABLE", it may be because your userid does not have the MOUNT attribute; go back to Step 1. If you have already done Step 1, then the problem is that the tape was mounted AFTER the MOUNT command was issued. Issue an UNLOAD console command and go back to step 1.

Now issue this TSO command to invoke the Tape Install program:

LOADGO 'FDR.INSTALL'

The Tape Install program (FDRLOAD) will be loaded from the tape and begin execution. Proceed to Step 4.

90.03 CONTINUED

STEP 3 Use this step to submit a batch job to copy the Tape Install program to a disk file, from which it can be executed under TSO.

Submit this jobstream:

```
//IEBGENER      EXEC      PGM=IEBGENER
//SYSPRINT      DD        SYSOUT=*
//SYSIN         DD        DUMMY
//SYSUT2        DD        DISP=(,CATLG),SPACE=(1600,50,RLSE),
//              DSN=user-specified-name,      <=== specify a dataset name
//              UNIT=SYSALLDA,VOL=SER=vvvvvvv <=== specify a disk volume
//SYSUT1        DD        DISP=OLD,DSN=FDR.INSTALL,LABEL=(9,EXPDT=98000),
//              UNIT=TAPE,                      <=== change if required
//              VOL=SER=FDR53T                  <=== change to FDR53P if production tape
```

After the successful completion of the IEBGENER job, issue this TSO command from ISPF Option 6 (TSO COMMANDS) or the TSO READY prompt:

LOADGO 'user-specified-name'

Specify the same data set name given in the JCL, in quotes.

The Tape Install program will be loaded from disk and begin execution. Proceed to Step 4.

STEP 4 The Tape Install program will prompt you for information on what, where and how to load the FDR tape files, in a series of four user-friendly screens. **No action will take place until you give the final confirmation on the fourth screen.** Only then are the output data sets allocated and cataloged with the names you specified, and the loading of those data sets begins (either in the foreground or via a batch jobstream).

Note: all dataset names and index name references are specified and displayed as fully-qualified names: a TSO userid will not be prefixed to the names unless you key it in.

SCREEN 1 - DATA SET SELECTION

```
WELCOME TO INNOVATION'S FDR TOTAL DASD MANAGEMENT SYSTEM INSTALLATION  SCREEN
1

PLEASE REPLY TO THE FOLLOWING PROMPTS. YOU WILL BE ABLE TO REVIEW AND
CHANGE YOUR SPECIFICATIONS PRIOR TO THE ACTUAL LOADING OF THE TAPE.

THE FOLLOWING DATA SETS MAY BE LOADED FROM THE INSTALLATION TAPE:

    1 - FDR INSTALLATION CONTROL LIBRARY
    2 - FDR LOAD MODULE LIBRARY
    3 - FDR ISPF DIALOG CLIST LIBRARY
    4 - FDR ISPF DIALOG PANEL LIBRARY
    5 - FDR ISPF DIALOG MESSAGES LIBRARY
    6 - FDR ISPF DIALOG SKELETON LIBRARY
    7 - FDR JCL LIBRARY

-----
---
<PRESS>  "ENTER"   -   SELECT ALL OF THE ABOVE DATA SETS AND CONTINUE
<TYPE>   "N,N,.. " -   SELECT THE SPECIFIED DATA SETS
<TYPE>   "END"    -   EXIT IMMEDIATELY
-----
---
                                PLEASE SELECT ONE OF THE OPTIONS LISTED ABOVE
SELECT ==>
```

This screen allows you to select which of the data sets are to be loaded from the FDR distribution tape. Normally, all data sets should be selected. When you are satisfied with the selection, press ENTER to continue to Screen 2.

90.03 CONTINUED

SCREEN 2 - DATA SET NAME SELECTION

```

----- DATA SET NAME SELECTION SCREEN ----- SCREEN 2

PLEASE REVIEW THE SELECTED DATA SET NAMES AND MAKE THE DESIRED MODIFICATIONS.

1 - INSTALL CONTROL..... IDP.ICLFDR53
2 - LOAD LIBRARY..... IDP.MODFDR53
3 - ISPF CLISTS..... IDP.DIALOG.CLIST
4 - ISPF PANELS..... IDP.DIALOG.PANELS
5 - ISPF MESSAGES..... IDP.DIALOG.MESSAGES
6 - ISPF SKELETON..... IDP.DIALOG.SKELETON
7 - JCL..... IDP.JCLFDR53

-----
<PRESS> "ENTER" - USE THE ABOVE SPECIFICATIONS AND CONTINUE
<TYPE> "ALL,INDEX" - ASSIGN NEW INDEX(ES) TO ALL DATA SET NAMES
<TYPE> "N,NEWNAME" - ASSIGN A NEWNAME TO THE DATA SET DESIGNATED BY "N"
<TYPE> "BACK" - GO BACK TO THE DATA SET SELECTION SCREEN 1
<TYPE> "END" - EXIT IMMEDIATELY
-----

PLEASE SELECT ONE OF THE OPTIONS LISTED ABOVE

SELECT ==>

```

This screen allows you to specify the data set names which will be used for the data sets you have selected to load from the tape. These may be existing data sets to be updated, or they may be new data sets which will be allocated and cataloged (new data sets are recommended).

The names shown above are the default names provided with the Tape Install program. You may change these names in one of 2 ways:

- 1) to change the current high-level index of all of the data sets to a different index (or indexes), enter "ALL,newindex(s)". For example,

ALL,FDR53 will change the names to FDR53.ICLFDR53, etc.

ALL,SYS3.FDR will change the names to SYS3.FDR.ICLFDR53, etc.

- 2) to completely change the name of any one data set, enter that data set's number followed by the replacement name. For example,

3,SYS2.IDP.LOAD will change the name of the load library.

You can use either or both of these techniques repeatedly until you are satisfied with the names.

If you intend to update an existing library, be sure that library name is correctly specified. However, we recommend that you always install into newly created libraries to avoid X37 ABENDs due to insufficient space in existing libraries. For new data sets, the install program will allocate them with sufficient space.

Note: during the installation, the ISPF dialogs will be modified to use the data set names you specify here. If you later rename the FDR libraries or copy the members to other libraries, you can update the dialogs with ISPF panel A.I.1A ([see Section 90.46](#))

Warning: if you are installing a trial tape (serial FDR53T) but you are already licensed for some components of FDR, be sure to install the tape into a set of libraries totally different from those used for the production components. All components on the trial tape, including those you are already licensed for, will be trial (date-protected) versions.

90.03 CONTINUED

SCREEN 3 - VOLUME SERIAL ISMS CLASS SELECTION

```

----- VOLUME SERIAL SELECTION SCREEN ----- SCREEN 3

THE FOLLOWING NEW DATA SETS WILL BE ALLOCATED AND CATALOGED:

      DISP  VOLUME  DATA SET NAME
1 - INSTALL CONTROL. NEW          IDP.ICLFDR53
2 - LOAD LIBRARY.... NEW          IDP.MODFDR53
3 - ISPF CLISTS.... NEW          IDP.DIALOG.CLIST
4 - ISPF PANELS.... NEW          IDP.DIALOG.PANELS
5 - ISPF MESSAGES... NEW          IDP.DIALOG.MESSAGES
6 - ISPF SKELETON... NEW          IDP.DIALOG.SKELETON
7 - JCL..... NEW          IDP.JCLFDR53

-----

<PRESS> "ENTER"      - USE THE ABOVE SPECIFICATIONS AND CONTINUE
<TYPE>  "ALL,VOLUME"  - ASSIGN A VOLUME TO ALL NEWLY ALLOCATED DATA SETS
<TYPE>  "N,VOLUME"   - ASSIGN A VOLUME TO THE DATA SET DESIGNATED BY "N"
<TYPE>  "SMS"        - DISPLAY SMS SPECIFICATIONS
<TYPE>  "BACK"       - GO BACK TO THE DATA SET NAME SELECTION SCREEN 2
<TYPE>  "END"        - EXIT IMMEDIATELY

-----

                        PLEASE SELECT ONE OF THE OPTIONS LISTED ABOVE

SELECT ==>

```

On this screen, the DISP column shows whether the install program found that the indicated data set already exists (OLD) or does not exist (NEW). For NEW data sets, you can specify volume and/or SMS information to be used for the allocation of the data sets. For OLD datasets, the VOLUME column shows the volume serial of the existing data set.

Similar to Screen 2, you can specify the target disk volume serial for all or any one of the data sets to be allocated. For example,

ALL,SYSLB2 will change the target volume serial for all the data sets.

2,SYSVOL will change the target volume serial for the load library.

You can use either or both of these techniques repeatedly until you are satisfied with the names. The volume serial can be omitted if the data set will be SMS-managed or if your system will allocate such data sets on non-specific storage volumes.

If SMS is active on your system, you may enter "SMS" which will take you to variations of Screen 3 which will allow you to specify the SMS storage class, management class, and/or data class to be assigned to each data set. However, this is not necessary if your installation's SMS ACS routines will assign proper classes to these data sets.

90.03 CONTINUED

SCREEN 4 - INSTALLATION PROCESSING OPTION

```

----- INSTALLATION PROCESSING OPTION SCREEN ----- SCREEN 4

PLEASE VERIFY THE FOLLOWING SPECIFICATIONS AND SELECT THE PROCESSING OPTION:

          DISP  VOLUME  DATA SET NAME
1 - INSTALL CONTROL. NEW  SYS032  IDP.ICLFDR53
2 - LOAD LIBRARY.... NEW  SYS032  IDP.MODFDR53
3 - ISPF CLISTS..... NEW  SYS032  IDP.DIALOG.CLIST
4 - ISPF PANELS..... NEW  SYS032  IDP.DIALOG.PANELS
5 - ISPF MESSAGES... NEW  SYS032  IDP.DIALOG. MESSAGES
6 - ISPF SKELETON... NEW  SYS032  IDP.DIALOG.SKELETON
6 - JCL..... NEW  SYS032  IDP.JCLFDR53

-----
<TYPE>  "FG"      -  START LOADING THE ABOVE DATA SETS IMMEDIATELY
<TYPE>  "BG"      -  CREATE THE JCL TO LOAD THE ABOVE DATA SETS
<TYPE>  "BACK"    -  GO BACK TO THE DATA SET SELECTION SCREEN 1
<TYPE>  "END"     -  EXIT IMMEDIATELY
-----

          PLEASE SELECT ONE OF THE OPTIONS LISTED ABOVE
SELECT ==>

```

On this screen, you can review all of the decisions you have made before starting the actual loading of the libraries from the distribution tape. Entering "BACK" on this screen (or any of the others) will allow you to go back and change options before installation.

If **FG** (foreground) installation is chosen, all the data sets indicated as NEW will be allocated, then IEBCOPY or other utilities are invoked under TSO to load each of the selected libraries from tape. This option appears only if you loaded the Tape Install program directly from tape (Step 2). FG is recommended since the tape is already mounted.

If **BG** (background) installation is chosen, the NEW data sets will be allocated under TSO (same as FG) but then batch JCL will be created to actually load the libraries from the tape. This JCL will be stored as member **FDRLOAD** in the ICL (Installation Control Library) you specified, or, if you did not select the ICL, in a dataset named "userid.FDRTEMP.JCL". **You must review this jobstream, make any changes necessary for your installation, and submit it for execution.**

If you can use ISPF, please proceed to [Section 90.10](#) after the tape is loaded.

If you don't have or can't use ISPF, please continue the installation process using the install instructions contained in [Section 91](#).

Notes:

- a. If IEBCOPY gives a non-zero return code, it is considered to be a serious error.
- b. If an ABEND Sx37 occurs, it is caused by lack of space in the disk data set. This should occur only when adding members to an existing data set since the libraries allocated by the Tape Install program should have sufficient space. Fix the offending data set either by compressing it, expanding its directory, allocating more space, moving to another volume or letting the Tape Install program create a new data set. LOGOFF and re-LOGON to free the existing allocations and restart the installation process.
- c. While loading the FDR CLIST library, the Tape Install program will update all references to the various FDR libraries to the names you have specified here. When you use the newly loaded FDR ISPF dialogs, all the proper library names will be preset for you. If you change ISPF library names after loading the tape, please [see Section 90.46](#) for instructions on updating the dialogs.

90.10 INVOKING THE ABR INSTALL ISPF DIALOG (PANEL A.I)

The FDR CLIST Library contains a member named ABRALLOC for installation and testing of the new FDR system. This CLIST invokes the FDR ISPF dialogs by concatenating the FDR ISPF libraries in front of your current ISPF library allocations. During the Tape Install process, the CLIST library was updated to reflect all of the FDR library names to which you loaded the new FDR version.

The following step (STEP 1) has to be repeated every time you logon to TSO until you permanently add the new FDR dialogs option to your installation's ISPF dialog as shown in [Section 90.46](#).

Note: All of the panels in this section and the following sections are identified by their ISPF "path" name, e.g., A.I.4. This assumes that "A" is the code used on your ISPF main menu to access the FDR dialog main menu; "A" is the code used on the temporary ISPF main menu used by the procedure below. When you install the FDR dialogs for production use, you may choose a different code (e.g., "F" which would change the panel names to "F.I.4", etc.).

- STEP 1** Issue the following command either under ISPF option 6 (TSO command processor), or under TSO "READY" mode:

EXEC 'fdr.clist.library(ABRALLOC)' <----	Use the data set name of the FDR CLIST Library that was specified in the installation process.
--	--

This command will allocate the FDR ISPF dialog libraries. If you issue the command under ISPF, skip step 2 and proceed to step 3.

- STEP 2** If you invoke the ABRALLOC CLIST from TSO "READY" mode, a standard ISPF primary menu displayed. Select (option "A") on this menu to access the FDR dialog.

ISPF PRIMARY OPTION MENU

Menu	Utilities	Compilers	Options	Status	Help

ISPF Primary Option Menu					
Option ==> A					
0	Settings	Terminal and user parameters		User ID . : DF	
1	View	Display source data or listings		Time. . . : 09:05	
2	Edit	Create or change source data		Terminal. : 3278	
3	Utilities	Perform utility functions		Screen. . : 1	
4	Foreground	Interactive language processing		Language. : ENGLISH	
5	Batch	Submit job for language processing		Appl ID . : ISR	
6	Command	Enter TSO or Workstation commands		TSO logon : V45ISPF	
7	Dialog Test	Perform dialog testing		TSO prefix: DF	
8	LM Facility	Library administrator functions		System ID : OS24	
9	IBM Products	IBM program development products		MVS acct. : **NONE**	
10	SCLM	SW Configuration Library Manager		Release . : ISPF 4.5	
11	Workplace	ISPF Object/Action Workplace			
A	FDR/ABR	FDR/ABR DASD Management Functions			
Enter X to Terminate using log/list defaults					

90.10 CONTINUED

STEP 3 Select option "I" (INSTALL) on the FDR Primary Options Menu to invoke the FDR Install dialog, as illustrated in the following figure:

PANEL A
FDR PRIMARY
OPTIONS
MENU

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- FDR PRIMARY OPTIONS MENU -----			
OPTION ==> I			
1	REPORTS	- ABR REPORTING FUNCTIONS	V 5.3/40
2	RESTOR	- ABR DATA SET RESTORE	
3	ARCHIVE	- ABR DATA SET ARCHIVE OR SUPERSCRATCH	
4	BACKUP	- ABR DATA SET BACKUP	
5	REMOTE Q	- ABR REMOTE QUEUE UTILITY FUNCTIONS	
C	COMPAKTOR	- COMPAKTOR MAP AND SIMULATION REPORTS	
R	RELEASE	- COMPAKTOR RELEASE	
I	INSTALL	- INSTALLATION AND MAINTENANCE OF FDR AND OPTIONAL PRODUCTS	
J	JCL PARMS	- SPECIFY FDR JCL AND SYSOUT DEFAULTS FOR SUBMITTED JOBS	
K	FORMAT	- MODIFY FORMAT OF GENERATED REPORTS	
M	MESSAGES	- FDR MESSAGES AND CODES QUERY FACILITY	
Q	QUERY	- FDR/ABR STATISTICS QUERY	
S	SRS	- SEARCH, REPORT, SERVICES DIALOG[LLS1]	
T	FDRTSEL	- BACKUP FILE MANAGEMENT UTILITY	

The FDR Install dialog provides for interactive specification of the Global Options for FDR, ABR, COMPAKTOR and FDRREORG and the dynamic creation of JCL streams necessary to complete the installation process.

RESTRICTING
THE USE OF THE
INSTALLATION
PANELS

If you want to be sure that only certain users can execute the FDR Installation dialogs, you can edit member ABRINPRI in the FDR Dialog CLIST library. ABRINPRI contains instructions for restricting user access.

To insure that users don't copy and modify ABRINPRI or other FDR CLISTs, you may be able to use your security system to grant execute-only access to the FDR Dialog CLIST library for unauthorized users.

90.11 SELECT THE FDR GLOBAL OPTIONS (PANEL A.I.4)

NEW INSTALLATIONS:

You will want to display each of the option panels that apply to the FDR components you are installing to review and set options appropriately. Innovation suggests that you review the text in the following sections of the manual as you go through this process, but you may also display the help panels for any option panel, with much of the same information.

EXISTING INSTALLATIONS:

You can copy the options and tables from your production version of FDR, but you may wish to review the option panels afterwards to be sure that any new options and options for new FDR components are properly set.

Many options which affect the FDR system may be changed permanently. Most such options are kept in a load module called FDROPT, in the FDR program library. Additional options for FDRREORG are stored in load module FDRREOOP. FDREPORT uses some options from FDROPT, but options unique to FDREPORT are stored in another options module (see [Section 54](#) for information on permanently changing FDREPORT options). Many of these options affect the defaults for various operands on FDR control statements and can be overridden at execution time; others may be changed only in FDROPT and cannot be overridden.

There are also some tables, such as the ABR lists, the COMPAKTOR unmovable table and the FDRREORG NOREORG table, which are also stored as load modules in the FDR program library.

There are two ways of modifying these options and tables. The FDR ISPF dialogs may be used to set all options and tables. This is the preferred way, since all options are displayed with online help to describe them. If you have not installed the dialogs or prefer not to use them, the options and tables may be set with program FDRZAPOP, the Global Option Change facility, described in [Section 91](#). Additionally, the FDRREORG options table and NOREORG table may be set with program FDRREOZO, described in [Section 31](#).

PANEL A.I
FDR
INSTALLATION
MENU

To enter the FDR dialog to display and set options in FDROPT, select option "4" (SETOPT) in the FDR Installation Options Menu, as illustrated in the following figure:

```

----- FDR TOTAL DASD MANAGEMENT SYSTEM -- INSTALLATION OPTIONS MENU -----
OPTION  ===>  4

      1  INSTALL  - LOAD SELECTED LIBRARIES FROM THE FDR DISTRIBUTION TAPE
      1A  DSNAMES - DISPLAY/CHANGE THE DATA SET NAMES OF THE FDR LIBRARIES
      2  LOADSAR  - LOAD THE STAND-ALONE PROGRAM (SAR) ONTO DISK

      4  SETOPT   - SET INSTALLATION OPTIONS IN THE FDR GLOBAL OPTIONS TABLE
      4A  DYNAM   - DISPLAY THE DYNAMICALLY INSTALLED FDR GLOBAL OPTIONS TABLE

      5  SETCPK   - SET UP THE COMPAKTOR UNMOVABLE TABLE
      5A  SETREORG - SET UP THE FDRREORG NOREORG LIST

      ABR OPTION INSTALLATION

      6  SETLIST  - SET UP THE ABR PROTECT LISTS AND RESTORE ALLOCATION LIST
      7  ABRCAT   - CREATE THE ABR CATALOG
      8  ABRVOL   - SET ABR DISK VOLUME PROCESSING OPTIONS
      9  BLDARC   - CREATE THE ARCHIVE CONTROL FILE
     10  DIALOG   - SET FDR DIALOG GLOBAL OPTIONS

     11  ADDISPF  - ADD FOR COMMANDS TO AN ISPF COMMAND TABLE

```

90.11 CONTINUED

**DYNAMICALLY
INSTALLED
OPTIONS**

If the Operating System exits are in use (as described in [Section 90.60](#)), the FDRSTART program will load a copy of the FDROPT option module into common memory (CSA). This in-memory copy of FDROPT will be used by the exits themselves and also by all jobs and TSO users using FDR programs **unless** they have a STEPLIB or JOBLIB pointing to an APF-authorized program library containing an FDROPT module. This allows you to give certain users special sets of FDR options.

Option 4A on the above menu can be used to display (but not modify) the options in effect in that in-memory module. The panels displayed by option 4A are identical to those described in the following sections except that the fields are not modifiable.

**FDR GLOBAL
OPTIONS**

The FDR Global Options Table (module FDROPT) contains installation options for security features, user exits, control statement defaults, etc. Options are organized by option type on the Global Options Primary Menu; each option on this panel takes you to another panel where you may display and modify the actual options.

**PANEL A.I.4
SET GLOBAL
OPTIONS**

```

----- FDR INSTALLATION -- SET FDR GLOBAL OPTIONS PRIMARY MENU -----
OPTION  ===>

      1 - SECURITY OPTIONS                      8 - MORE ABR GENERAL OPTIONS
      2 - GENERAL OPTIONS                      9 - ABR REPORT DEFAULTS
      3 - COMPAKTOR OPTIONS                    10 - MORE ABR REPORT DEFAULTS
      4 - ABR GENERAL OPTIONS                  11 - OPERATING SYSTEM EXITS
      5 - ABR DATA SET NAMES                  12 - FDRREORG OPTIONS
      6 - ABR ARCHIVE UTILITY DEFAULTS         13 - FDRCLONE OPTIONS
      7 - ABR DISK PROCESSING OPTIONS

      SAVE - SAVE OPTION CHANGES              COPY - COPY OPTIONS FROM A PRIOR LEVEL
      CANCEL - EXIT WITHOUT SAVING CHANGES    AUDIT - DISPLAY USER CHANGED OPTIONS
      REFRESH - REFRESH OPTIONS TABLE IN LPA  RESET - RE-INITIALIZE ALL OPTIONS

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME===> 'IDP.MODFDR53'
  VOLUME SERIAL===>

NOTE: TO REFRESH THE OPTIONS THAT ARE DYNAMICALLY INSTALLED IN THE ACTIVE LPA,
IT IS NECESSARY TO RUN FDRSTART - USE THE REFRESH CMD TO GENERATE FDRSTART JCL.

```

The name of the FDR program library used during installation will be displayed. If necessary, correct that library name; you can also specify the volume serial of the library if it is not cataloged. These values will be saved in your ISPF profile so that you do not need to reenter them in the future. The subsequent panels will display the options currently in effect in the FDROPT module in that library, and will update that library when the options are saved.

90.11 CONTINUED

**SET OPTION
COMMANDS**

As shown, several special commands are available on this panel:

- **SAVE** updates the FDROPT module in the specified program library. No changes are made to the library until SAVE is entered, so you may freely switch between option panels and change options until you are satisfied.
- **CANCEL** exits to the previous menu and discards all options changed since the last SAVE command.
- **REFRESH** can be used to dynamically refresh the in-memory copy of FDROPT (see "Dynamically Installed Options" on the previous page) or to install the FDR Operating System exits if not already active. It can also refresh the LLA copy of FDR tables and lists after they are updated.
- **COPY** is used to copy options and tables from a previous release of FDR (see next page).
- **AUDIT** will create and display a data set containing the values and descriptions of all options which are not currently set to the Innovation default, i.e., an audit of all changed options. You can print this data set.
- **RESET** resets all options to the default values distributed by Innovation.

Warning: RESET should be used with care since it does reset all options to their defaults. This is especially dangerous for existing customers, since this may change the operation of FDR. We suggest that you use the AUDIT function (above) to document options not set to the default before using RESET.

SAVE and CANCEL can be used on any of the option panels. Other commands are only valid on this panel. If an option value is changed since the last SAVE command and you attempt to exit from this panel, you will be prompted to SAVE or CANCEL the changes before exiting.

NEW INSTALLATIONS

Innovation suggests that **new installations** should display each option panel and review the options on them. However, the FDR Global Options that are especially important for **new installations** are:

For all installations:

- Security Options (panel A.I.4.1, [Section 90.12](#))

For ABR installations:

- ABR Data Set Names (panel A.I.4.5, [Section 90.16](#)).
- If you have a tape management system of any kind set the TMS option (panel A.I.4.4, [Section 90.15](#)).
- If you have an automated tape library (usually called an ATL or silo), you may need to set the DYNDEALC option (panel A.I.4.4, [Section 90.15](#)).

90.11 CONTINUED

EXISTING INSTALLATIONS

The COPY command on the FDR Global Options Panel (A.I.4) can be used to copy the option values set in a previous version of FDR by reading the FDROPT and FDREOOP module from the previous load library and setting the equivalent options in the new FDROPT and FDREOOP . At the same time, you can also copy the ABR Restore Allocation List, the ABR Protect Lists, the COMPAKTOR Unmovable Table and the FDRREORG NOREORG Table, as illustrated in the following figure.

Warning: Do not copy the FDROPT or FDRREOOP modules from a previous version of FDR using any standard copy utility (such as IEBCOPY or ISPF COPY); the dialog COPY function copies option values while preserving version information and new defaults in the new FDROPT.

However, you can use IEBCOPY or ISPF COPY to copy the table/list modules from a previous version of FDR. In fact, you must do so if you have expanded any of those table/list modules so that the expanded size will be retained in the new release.

```

----- FDR INSTALLATION -- COPY FDR GLOBAL OPTIONS -----
COMMAND ===>

COPY "TO" DATA SET:  IDP.MODFDR53

TO CANCEL THE COPY OPERATION, EITHER PRESS THE END KEY (PF3) OR TYPE "CANCEL".

SPECIFY "FROM" DATA SET BELOW.

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME===> 'FDR.PROD.LOAD'
  VOLUME SERIAL===>

IN ADDITION, COPY THE FOLLOWING OPTIONS MODULES:

  ALLOCATE - ABR RESTORE ALLOCATION LIST... ==> NO      (YES NO)
  ARCPROT  - ABR ARCHIVE PROTECT LIST..... ==> NO      (YES NO)
  ABRPROT   - ABR BACKUP  PROTECT LIST..... ==> NO      (YES NO)
  RESTPROT  - ABR RESTORE PROTECT LIST..... ==> NO      (YES NO)
  SCRPROT   - ABR SCRATCH PROTECT LIST..... ==> NO      (YES NO)
  CPKUNMOV  - COMPAKTOR UNMOVABLE TABLE..... ==> NO      (YES NO)
  FDRNORG   - FDRREORG NOREORG LIST..... ==> NO      (YES NO)

```

Enter the data set name (and optional volume serial) of the FDR program library containing the previous version of FDR. Only those options which are not set to the Innovation default in effect for that version are copied; this way, if the Innovation default for an option is changed in the new version, the new default will not be overridden with the old default. The option values copied are immediately saved in the "to" program library; no SAVE command is required.

To copy one or more of the lists and tables at the same time, enter YES next to each selected list/table. The copied tables and lists are immediately saved in the "to" program library.

90.12 SECURITY OPTIONS (PANEL A.I.4.1)

The programs in the FDR system, as a default, will not invoke any type of security when dumping or restoring data sets or volumes. Since FDR does not open individual datasets, most security checks will be bypassed for FDR operations unless you enable the security option documented below. **By default, all security options are disabled and no security checking is done.**

SECURITY RECOMMEND- ATIONS

All installations with a SAF-compatible security system, including RACF, Top Secret and ACF2, should enable the ALLCALL option to activate security checking. ACF2 users should consult their security documentation; prior to V6 of ACF2, customization was required to properly handle SAF (RACROUTE).

WARNING: the FDR security exits FDRYOPEN and FDRYPASS, formerly provided by ACF2, are no longer used for FDR security; if currently in use, you should disable them and convert to ALLCALL.

PANEL A.I.4.1 SECURITY OPTIONS

```

----- FDR INSTALLATION -- SET FDR GLOBAL SECURITY OPTIONS -----
COMMAND ==>

ALLCALL  RACF ALWAYS CALL OPTION ENABLED . . . . . NO

NOABSTRK ABSOLUTE TRACK OPERATIONS ALLOWED . . . . . YES

NONEW    RENAME USING NEWDD, NEWNAME, NEWINDEX AND NEWGROUP ALLOWED . . . . YES

```

ALLCALL

If set to YES, FDR does SAF-compatible security checks for volumes and data sets. ALLCALL is set to NO (disabled) by default.

ALLCALL causes a SAF call in the form of RACROUTE REQUEST=AUTH to be used for both volume-level protection and data set protection. The level of authority required (READ, UPDATE, or ALTER) depends on the function being executed. The types of security calls issued by ALLCALL for each type of FDR operation are detailed in [Section 80.15](#).

NOABSTRK

If set to NO, prevents FDR from doing any absolute track operations (SELECT FROM/TO) since no dataset-level checking can be done for this type of disk access. By default, it is set to YES (absolute track operations allowed).

Program FDRDSF can DUMP, RESTORE or PRINT physical disk tracks, and program FDRCOPY will COPY physical disk tracks, if the user specifies the absolute track operands (FROM and TO). FDRDSF and FDRCOPY do not attempt to determine the data set(s) in which the requested tracks reside, so security checking for data set names is not done.

If ALLCALL is set to YES, security checking for absolute track operations is done at the volume level (see "DASDVOL Volume Protection" in [Section 80.15](#). If a DASDVOL profile exists for the volume being accessed, the user must have READ authority for DUMP or PRINT or for the input volume on COPY, or ALTER authority for RESTORE or for the output volume on COPY. But if no DASDVOL profile exists, then **all** users will be able to do absolute track operations on the volume unless NOABSTRK is set to NO.

NONEW

If set to NO, datasets may not be restored or copied to a new name. This will cause the operands NEWNAME, NEWGROUP, NEWDD and NEWINDEX to be treated as invalid. By default it is set to YES (restore to new name allowed).

90.13 GENERAL OPTIONS (PANEL A.I.4.2)

These options are used by most FDR programs. They can be displayed or changed on the following panel.

**PANEL A.I.4.2
GENERAL
OPTIONS**

----- FDR INSTALLATION -- SET FDR GLOBAL GENERAL OPTIONS ----- COMMAND ==>	
ICFCORE	ICF VSAM CLUSTER AND COMPONENT NAMES TABLE SIZE (BYTES) .50000
FDRCC	FDR AND FDRDSF ERROR RETURN CODEABEND
FDRSTMT	PROGRAM FDR MUST REQUIRE PARM OR SYSIN INPUTNO
LINECNT	MAXIMUM NUMBER OF LINES TO BE PRINTED ON REPORTS58
SELTERR	TREAT DATA SET SELECTION FAILURE AS AN ERRORYES
ALCRSTIN	INTERFACE WITH ALLOCATION CONTROL PRODUCT FOR NON-VSAM ..NO
RESTLRDT	ALWAYS RESTORE DATA SET LAST REFERENCE DATE FROM BACKUP .NO
RESTRCDT	ALWAYS RESTORE DATA SET CREATION DATE FROM BACKUPNO
RESTEXTD	ALWAYS RESTORE DATA SET EXPIRATION DATE FROM BACKUPNO
HFSQUIESCE QUIESCE HFS DURING BACKUP..... NO	
ROUTECD	WTO ROUTING CODES(2,11)
DESCRPTCD	WTO DESCRIPTOR CODES(2)

- ICFCORE** In order to match ICF VSAM component names to cluster names, the FDR dataset dump programs must build a table of such names. The default size of this table is 56K if the input disk is a 3390 (which holds about 650 components) and 48K for other disks (about 600 components). If any input disk contains more than this number of components, the table must be enlarged. The default value for ICFCORE (50000) causes the above defaults to be used. Setting ICFCORE to any other value, in bytes, will cause a larger table to be used and will increase the region required per TAPEX DD statement by that amount (the default values do not add any extra storage). If the ICFCORE value is not high enough, FDR will issue message FDR152 REASON=9.
- FDRCC** If an error occurs during execution of FDR, FDRDSF, FDRCOPY, or FDRTCOPY that is not severe enough to cause immediate termination, these programs complete their processing and then issue an ABEND or set a return code to call attention to the error. The default (ABEND) causes an ABEND U0888. If you wish a return code instead of the ABEND, set any desired value from 8 to 255.
- FDRSTMT** If set to NO and program FDR is executed with no control statements in SYSIN and no PARM= on the EXEC statement, it will default to a dump (backup). However, if the intention was to do a restore but the control statement or parm was accidentally omitted, this may result in overwriting the backup tape from which you intended to restore. Changing this option to YES will cause FDR to fail if no statement or parm was specified, requiring that the operation to be performed be positively specified. The default is NO.
- LINECNT** You can override the maximum number of lines to be printed on any report page produced by the FDR system to any value from 10 to 99. The default is 58 lines per page. This option can be overridden at execution time.
- SELTERR** The FDR system will produce a diagnostic message if a SELECT or EXCLUDE or related control statement, specifying a data set name or filter, does not actually match any of the data sets processed in the step (in other words, the control statement was not used). If set to YES, FDR considers this a possible user error and causes a return code or U0888 ABEND at step end to draw attention to it. If you do not want to consider this an error condition, change SELTERR to NO (the diagnostic message will still be printed but it will not be considered an error). This option can be overridden at execution time. The default is YES.

90.13 CONTINUED

ALCRSTIN If set to YES, enables support for the interface to several allocation control software products. The interface invokes a disk allocation management software product to get target volume serials when FDR is allocating a non-VSAM data set. The interface is supported by several disk allocation products:

- **POOL-DASD** from Boole and Babbage
- **PROSMS** V3.5 (or higher), with the "FDR Assist" feature enabled, from Boole and Babbage
- **ACC** from DTS Software
- **SAMS:ALLOCATE** from Sterling Software (check with Sterling for release information)

This support requires a interface module, PLDRSTIN, which is provided by each of the vendors of disk allocation products listed above. This module must be located in some linklist library so that it is available to FDR whenever data sets are restored. If this option is enabled and PLDRSTIN is available, FDR will call PLDRSTIN to obtain target volumes for non-VSAM data sets. The default is NO.

Note that support for all of the above disk allocation products when allocating VSAM clusters is automatic and is used even if the ALCRSTIN option is disabled. For SMS-managed data sets, SMS may override any volume decisions made by the allocation products.

If you have any questions or problems installing the PLDRSTIN module from your allocation product vendor, please check first with that vendor for instructions and additional information. If you can't resolve the problem, call Innovation for assistance.

In previous releases of FDR, this was known as the POOLDASD option, after the product for which the interface was first designed. To avoid confusion it has been renamed but it's operation is the same.

RESTLRDT
RESTRCDT
RESTEXDT During a data set restore, copy or move, the three date fields in the Format-1 DSCB of the data set (last reference date, creation date, and expiration date) are normally set based on rules described in [Section 80](#). Sometimes the date fields are copied from the backup or input data set; sometimes they are changed to today's date or zero. For each of those three dates, a NO value in the associated option will cause those rules to be used.

If you change any or all of these options to YES, the associated date field will always be copied from the value on the backup tape or the value from the input data set.

Note: In V5.2 of FDR, these options were implemented as custom zaps (C-52.0012 and C-52.0014). These zaps are no longer required in V5.3; use these options instead.

WARNING: if you are archiving data sets based on their last reference date (the **ADAYS=** operand or similar values in SMS management classes), you should not set **RESTLRDT** to YES since this restores the last reference date that the data set had when it was archived and may make it immediately eligible for archive again.

HFSQUIESCE If set to YES, during all types of backups, FDR will attempt to quiesce access to all HFS data sets (Hierarchical File Systems, used by Unix System Services for OS/390). This is required to get a consistent backup of the HFS data if updates may be occurring during the backup.

ROUTE CODE/
DESCRIPTOR
CODE The WTO routing and descriptor codes used for most of the operator messages in the FDR system can be changed. Multiple codes can be specified for each. The default values for the routing codes are 2 (Operator Information) and 11 (Programmer Information) and the default descriptor code is 2 (Immediate Action Required). The routing and descriptor codes are documented in the IBM manual "Routing and Descriptor Codes" and are used by MVS to control on which consoles the messages will appear, and how they will be displayed.

90.14 COMPAKTOR OPTIONS (PANEL A.I.4.3)

The options on this panel apply only to COMPAKTOR (program FDRCPK). Note that COMPAKTOR is also affected by some of the security options (panel A.I.4.1) and some of the FDR general options (panel A.I.4.2) described in preceding sections. **If you are not licensed for COMPAKTOR you may skip this panel.**

**PANEL A.I.4.3
COMPAKTOR
OPTIONS**

```

----- FDR INSTALLATION -- SET FDR GLOBAL COMPAKTOR OPTIONS -----
COMMAND ===>

ACTMESS  ISSUE MESSAGE FDRW81 'DETECTS VOLUME/UNIT IN USE' ..... YES

CPKCC    COMPAKTOR ERROR RETURN CODE ..... ABEND

CPKABSUN TREAT ALL ABSOLUTE TRACK DATA SETS AS UNMOVABLE ..... NO

HILIGHT  USE OVERPRINTING IN MAP REPORTS FOR FIELD CONTRAST ..... YES

NOSECOND ALLOW RELEASE IF NO SECONDARY ALLOCATION IS DEFINED ..... YES

SIZEKEEP CONTIGUOUS EXTENT SIZE TO TRY TO KEEP IN PLACE (TRACKS) ..... 100
FREEAREA TARGET PERCENT REDUCTION IN THE NUMBER OF FREE SPACE AREAS ..... 90
MAXKEEP  MAX NUMBER OF DATA SET EXTENTS TO BE ELIGIBLE FOR SIZEKEEP ..... 60

```

ACTMESS

If set to YES during an old-style COMPAKTion from a backup (COMPAKT TYPE=COMPAKT), if any other tasks have the volume allocated COMPAKTOR will issue message FDRW81 to warn the operator. This message requires a reply to continue or bypass the COMPAKTion. If all of the precautions listed in [Section 40.21](#) ("COMPAKTing Active Volumes") are observed for all your COMPAKTOR jobs, then it is safe to COMPAKT a volume while it is in use, and this message is not needed.

You can permanently suppress the message for all COMPAKTOR jobs by changing the option to NO on this panel. You can also override it for any COMPAKTOR job by specifying ACTMESS=NO on the COMPAKT statement (or specify ACTMESS=YES to override NO on this panel). The default is YES.

ACTMESS does not apply for FAST COMPAKTion (COMPAKT TYPE=FCPK) or RELEASE (COMPAKT TYPE=RLSE).

CPKCC

If an error occurs during execution of COMPAKTOR that is not severe enough to cause immediate termination, it completes processing of the current volume (and any other requested volumes) and then issues an ABEND or sets a return code to call attention to the error. The default (ABEND) causes an ABEND U0888 at step termination. If you wish a return code instead of the ABEND, change CPKCC to any desired value from 8 to 255. However, if CPKCC is set to ABEND but FDRCC (panel A.I.4.2, [Section 90.13](#)) is set to a return code, COMPAKTOR will use the FDRCC return code.

Severe errors will cause immediate termination of the COMPAKTOR step with one of several other ABEND codes.

CPKABSUN

If set to NO, COMPAKTOR will treat data sets which were allocated by absolute track allocation (SPACE=(ABSTRK,...)) as movable. If you want COMPAKTOR to treat them as unmovable and always leave them in their current locations, change this option to YES. The default is NO.

HILIGHT

If set to YES, COMPAKTOR will highlight certain data on its volume maps in order to make it stand out, e.g., free space areas. This is done by overprinting the line three times. This is valuable only on impact printers (which actually strike the paper three times). On most laser printers the overprint is not noticeable, and when viewed on TSO terminals usually results in displaying the line 3 times. If your COMPAKTOR output is usually printed on laser printers or viewed on TSO, you can suppress the overprinting by changing this option to NO. The default can be overridden at execution time by specifying HILIGHT=YES or NO. The default is YES.

90.14 CONTINUED

NOSECOND The COMPAKTOR function of releasing excess space, during either a TYPE=RLSE step or a full COMPAKTion, may be limited to data sets that have a secondary allocation value defined; this allows you to exclude data sets that have only primary allocation defined, and are unable to acquire secondary extents if they need to re-expand. If set to YES, secondary allocation is not considered when releasing space. To change the default so that COMPAKTOR will not release space from data sets that do not have a secondary allocation value, change this option to NO. The default can be overridden at execution time by specifying NOSECOND=NORLSE or NOSECOND=RLSE. The default is YES.

Innovation recommends that NOSECOND be set to NO so that space release will bypass data sets created without secondary allocation.

SIZEKEEP The SIZEKEEP operand of COMPAKTOR, explained in [Sections 40.08](#) and [40.14](#), has 3 suboperands. The default values for those suboperands are specified by the options SIZEKEEP, FREEAREA, and MAXKEEP respectively. The distributed defaults are SIZEKEEP=100 (tracks), FREEAREA=90 (percent), and MAXKEEP=60 (extents). You can change the defaults for any or all of these values. However, the defaults can also be overridden at run time by the SIZEKEEP operand on the COMPAKT statement.

Innovation recommends that you not change the default values for SIZEKEEP (100,90,60) .

90.15 ABR GENERAL OPTIONS (PANEL A.I.4.4)

These options are used by components of ABR. Note that ABR is also affected by some of the security options (panel A.I.4.1) and some of the FDR general options (panel A.I.4.2) described in preceding sections, and there are additional option panels (described in succeeding sections) which also affect ABR. **If you are not licensed for ABR you may skip this panel.**

**PANEL A.I.4.4
ABR GENERAL
OPTIONS**

----- FDR INSTALLATION -- SET FDR GLOBAL OPTIONS FOR ABR -----		
COMMAND ==>		
ABRCC	ABR ERROR RETURN CODE	12
DISKUPDATE	SET DSCB FLAGS FOR REMOTE QUEUE BACKUP AND ARCHIVE DUMPS	YES
TMS	TAPE MANAGEMENT SYSTEM PRODUCT IS INSTALLED ON THIS SYSTEM	NO
ARCRCALL	AS A DEFAULT, ARCHIVE WILL RECATALOG FOR AUTO-RECALL	NO
MIGRAT	USE VOLSER "MIGRAT" ON ARCHIVE RECATALOG FOR AUTO-RECALL	NO
ARCCOPY	DEFAULT ARCHIVE BACKUP COPY NUMBER FOR ARCHIVE RESTORES	1
BKPCOPY	DEFAULT BACKUP COPY NUMBER FOR BACKUP RESTORES	1
DYNDEALC	FOR DYN TAPE, REALLOCATE THE TAPE UNIT FOR EACH VOLUME SET	NO
ALLOCATEFLAG - DYNAMIC ALLOCATION FLAGS		
WTVOL	WAIT FOR VOLUMES	NO
WTDSN	WAIT FOR DATA SETS (IF JES3 THEN SPECIFY 'NO')	NO
WTUNT	WAIT FOR UNITS	NO
OFFLN	CONSIDER OFFLINE VOLUMES	NO
MOUNT	VOLUME MOUNT PERMITTED	NO

ABRCC

If an error occurs during execution of ABR that is not severe enough to cause immediate termination, ABR will complete processing of the current volume and other requested volumes and then terminates the step with an ABEND or return code to call attention to the error. The default of 12 sets return code 12. If you wish to change the return code, set any desired value from 8 to 255. Set ABRCC to "ABEND" if you wish ABR to end with ABEND U0888.

DISKUPDATE

If set to YES, the ABR remote queue utility (FDRABRUT) will update the DSCBs of selected data sets directly when a remote queue backup or archive is requested. FDRABRUT must be executing as an authorized program to do so. If FDRABRUT may be executed unauthorized, or if you don't want to update the DSCBs directly, change this option to NO; this will cause FDRABRUT to use remote queue datasets (see "Remote Queue Datasets" later in this section and sections [50.03](#) and [51.03](#)). FDRABRUT will always use remote queue datasets for restores from backup and archive. The default is YES.

Note that when FDRABRUT is executed because of a request for a remote queue operation from the FDR ISPF panels (the most common case), there is a dialog option which specifies whether to force DISKUPDATE=YES or NO or to use the default in the Global Option Table. As distributed, it uses the global value, but this can be changed on panel A.I.10.

TMS

The tape techniques used by ABR are generally compatible with tape management systems, such as CA-1, CA-TLMS, IBM's DFSMSrmm, etc. However, you must inform ABR that Tape Management exists in the Operating System by specifying YES. Note that the TMS option refers to ALL tape management systems, not just CA-1 (formerly known as TMS). Although most ABR functions will work successfully with a TMS even with this option disabled (NO), the ABR LASTAPE function will fail under a TMS without the TMS option. For details on "How ABR Interfaces with Tape Management Systems", see [Section 80](#). The default is NO.

90.15 CONTINUED

- ARCRCALL** If set to NO, ABR ARCHIVE will not catalog archived data sets for automatic recall unless the RECALL or RECALL=YES options are specified on the DUMP TYPE=ARC statement. Since many installations want most archived datasets to be eligible for auto-recall, you can make RECALL=YES the default by changing this option to YES. This default can still be overridden for specific ARCHIVE runs with RECALL=NO. This option has no effect on ABR application backup (DUMP TYPE=APPL) and when archiving data sets which are not already cataloged.
- Recommendation: Set ARCRCALL to YES unless you do not intend to use auto-recall.**
- MIGRAT** By default, if ABR ARCHIVE recatalogs archived data sets for automatic recall (RECALL=YES in effect), the catalog entry will retain the original volume serial of the disk from which the data set was archived, unless the MIGRAT=YES option is specified on the DUMP TYPE=ARC statement, in which case the volume serial "MIGRAT" is used instead. MIGRAT=YES is required if you archive SMS managed data sets or are running ISPF V4 or above but it is recommended for ALL installations. You can make MIGRAT=YES the default by changing this option to YES. This default can be overridden for specific ARCHIVE runs with MIGRAT=NO. The default is NO.
- Recommendation: Set MIGRAT to YES.**
- ARCCOPY** If set to "1", ABR ARCHIVE restores and FDRAPPL application restores will restore from the COPY1 backup (unless COPY1 has reached its expiration date, in which case COPY2 will be used if it exists). You can make COPY2 the default by changing this option to "2". However, if COPY2 is expired or unavailable, ABR will NOT switch back to COPY1. This may be useful at a disaster site where only COPY2 is available from offsite storage. This default can be overridden for any given restore with the COPY=n option on the RESTORE or SELECT statement. The default is "1".
- BKPCOPY** If set to "1", ABR restores from full-volume and incremental backups will restore from the COPY1 backup (unless COPY1 is no longer cataloged, in which case COPY2 will be used if it exists). You can make any copy ("1" through "9") the default by changing this option. However, if the specified copy is no longer cataloged, ABR will NOT switch back to COPY1. This may be useful at a disaster site where only COPYn is available from offsite storage. This default can be overridden for any given restore with the COPY=n option on the RESTORE or SELECT statement. The default is "1".
- DYNDEALC** Controls the operation of the DYNTAPE option for ABR restores from tape.
- If set to NO, ABR will allocate a tape drive for the first tape required for a given restore step, but if the next tape required is the same type as the previous tape (e.g., 3490E) the same drive will be used to mount that tape. Only when a different type of drive is required for the next tape will ABR deallocate the current drive and allocate a new one. This is to avoid the problem of the ABR restore tapes bouncing from one drive to another and possibly waiting for a free drive. But, in installations with an automated tape library (often called an ATL or silo) specific tape volumes may need to be allocated to specific drives so that the ATL can service them; if the next tape ABR requires is in a different ATL or not in an ATL at all, the mount request may require manual operator action or may fail.
- If set to YES, ABR will deallocate and reallocate whenever a new tape backup data set on a new volume is required, allowing the allocation to be directed to a tape drive appropriate for the location of the tape, even if it is in a different ATL or outside an ATL. If the next backup file is on the tape already mounted, it will not be reallocated. **All users with ATLs or silos should enable this option. The default is NO.**
- When restoring from disk backups, ABR will always reallocate each backup dataset.

90.15 CONTINUED

**ALLOCATEFLAG
-DYNAMIC
ALLOCATION
FLAGS**

ABR will dynamically allocate the backup medium (tape or disk) during restore operations if the DYNTAPE operand is specified (assumed for auto-recall). The options below affect dynamic allocation of the backup data sets.

MOUNT	If set to YES, allows a TSO user to mount a tape, whether or not the user has MOUNT authority in their TSO profile. This can be used to allow TSO users to do foreground recalls from tape even though they are not normally allowed to mount tapes for their TSO session.
OFFLN	If set to YES, allows dynamic allocation to ask the operator to vary an offline tape drive online to mount the backup tape if no online tape units of the proper type are available. For batch jobs and TSO users with MOUNT authority in their TSO profile, offline units are considered by MVS even if the OFFLN option is not specified so you usually do not have to enable it. However, if the MOUNT option above is specified, then offline units will not be considered for any ABR allocations, so if either MOUNT or OFFLN is enabled, the other should also be enabled.
WTDSN	If set to YES, allows ABR to wait if the required backup data set name is in use. This could be true if another restore is currently using it, or if the backup/archive job creating the backup data set is still running. Allocation will automatically proceed when the data set is released by the other job. This option is recommended for systems with JES2; however, it should not be set in systems with JES3, where it can cause an interlock.
WTUNT	If set to YES, allows ABR to wait for an available unit if all eligible units are already in use. The operator may need to reply WAIT to an allocation recovery message unless you have automated this reply. This option is recommended for all systems.
WTVOL	If set to YES, allows ABR to wait for the tape volume if it is currently in use by another job. This could be true if another restore is currently using it, or if the backup/archive job is still running. WTVOL is automatically assumed on full volume restores. This option is recommended for all systems, particularly for installations doing auto-recall since multiple recalls may require the same tape volumes.

NOTE: The WTDSN, WTUNT, and WTVOL options can in some circumstances result in ABR being in a WAIT state for a very long time. If this wait exceeds your installation WAIT limit (typically 30 minutes) ABR may ABEND with a S522 (excessive WAIT time). If this occurs you may have to specify TIME=1440 on the ABR EXEC statement to suppress this.

Recommendation: Activate the options WTDSN, WTUNT, and WTVOL (except do not activate WTDSN under JES3). These options are not enabled (set to NO) by default.

90.16 ABR DATA SET NAMES (PANEL A.I.4.5)

These options control naming conventions and data set names used by ABR. Quotes are not required around any of the names specified on this panel; they are assumed to be fully-qualified names. **If you are not licensed for ABR you may skip this panel.**

**PANEL A.I.4.5
ABR DATA SET
NAMES**

```

----- FDR INSTALLATION -- SET FDR GLOBAL DATA SET NAME OPTIONS -----
COMMAND ===>

ABRINDEX HIGH LEVEL INDEX FOR BACKUPS AND ABR MODEL DSCBS..... FDRABR.
SCRINDEX HIGH LEVEL INDEX FOR DATA SETS IN THE ABR SCRATCH CATALOG.. #.

DSNCK      ARCHIVE AND REMOTE QUEUE DATA SET NAMES ARE STANDARD..... YES

LASTAPE    LAST TAPE OPTION DATA SET NAME PREFIX..... FDRABR.LASTAPE

POOLDISK   POOLDISK  OPTION DATA SET NAME PREFIX..... FDRABR.POOLDISK

ARCDN      ARCHIVE CONTROL FILE.... FDRABR.ARCHIVE

          REMOTE QUEUE DATA SETS
ARCHDMPQ   ARCHIVE DUMP QUEUE..... FDRABR.ABRARDQ.DATA
ARCHRSTQ   ARCHIVE RESTORE QUEUE... FDRABR.ABRARCH.DATA
BKUPDMPQ   BACKUP DUMP QUEUE..... FDRABR.ABRBKDQ.DATA
BKUPRSTQ   BACKUP RESTORE QUEUE.... FDRABR.ABRREST.DATA

```

ABRINDEX ABR normally uses a first level index of 'FDRABR' for all backup and ARCHIVE data sets created by FDRABR ([Sections 50](#) and [51](#) describe the ABR naming conventions). In addition, this index level is used in the names of the ABR MODEL DSCBs that are placed in the VTOCs of every volume initialized for ABR processing ("FDRABR.Vvolser"). It is highly recommended that the standard index name of 'FDRABR' be used but if you must use a different name change it here; ABRINDEX must be a single index (ABR will add a period at the end if you omit it).

WARNING: once you have begun to use ABR in production, you MUST NOT change the value of ABRINDEX. Because of the way that backup file names are stored in ABR, changing ABRINDEX will cause ABR to be unable to locate any existing backup or archive data sets. If you must change it, the change must be made before any production ABR runs are made. Whatever index you choose, it must be used only for ABR data sets; do not choose an index that will be used for any other data sets.

SCRINDEX This single character (alphabetic or national) plus a period will be prefixed onto the name of any data set that is scratched by MVS, if it also has a current ABR backup. This modified name will be cataloged into the ABR scratch catalog if the ABR DADSM Pre-processing exit is installed (see [Section 90.60](#)). The default is "#." (pound-sign) and Innovation recommends that you do not change it but if you must change it specify a single replacement character here (ABR will add a period at the end if you omit it).

DSNCK If set to YES, ABR validates the names of the Archive Control File and the remote queue datasets, as explained under those options. To disable this validation, change this option to NO. The default is YES.

LASTAPE This value is the prefix for the data set name used for the LASTAPE option; LASTAPE allows multiple ABR jobs to place backup files on the same tapes as explained in [Sections 50.04](#) and [51.04](#). LASTAPE must itself be a valid MVS data set name; it can be up to 16 characters long and can consist of any number of index levels. The default is 'FDRABR.LASTAPE'.

Note: If you have a tape management system, you **must** also enable the TMS option on Panel A.I.4.4 if you intend to use the LASTAPE feature.

90.16 CONTINUED

- POOLDISK** This value is the prefix for the data set name used for the POOLDISK option; POOLDISK allows ABR to manage a pool of disks used for ABR output files (usually used for ARCHIVE) as explained in [Section 51.04](#). POOLDISK must itself be a valid MVS data set name; it can be up to 16 characters long and can consist of any number of index levels. The default is 'FDRABR.POOLDISK'.
- ARCDSN** This is the fully-qualified name of the default ARCHIVE Control File (up to 44 characters). ABR will dynamically allocate this file for the AUTOMATIC RECALL of data sets and for any ARCHIVE job specifying the DYNARC option. Note that you may have many ARCHIVE Control Files but only the one named here can be used for auto-recall. One index level of the name (other than the first) must be 'ARCHIVE' unless the DSNCK option has been set to NO. The default name is 'FDRABR.ARCHIVE'; if you prefer to use a different name, change it here. The ARCHIVE Control File can be created on panel A.I.9 ([see Section 90.43](#)).
- ARCHDMPQ** This is the fully-qualified name of the default ARCHIVE DUMP Remote Queue data set (up to 44 characters). The Remote Queue utility, FDRABRUT, will dynamically allocate this data set and add ARCHIVE DUMP requests to it, if the DISKUPDATE option is set to NO (if DISKUPDATE is YES, this data set is not used). One index level of the name (other than the first) must be 'ABRARDQ' unless the DSNCK option has been set to NO. The default name is 'FDRABR.ABRARDQ.DATA'; if you prefer to use a different name, change it here. [See Sections 51.04](#) and [90.45](#) for an explanation of the use of this data set by ABR.
- ARCHRSTQ** This is the fully-qualified name of the default ARCHIVE RESTORE Remote Queue data set (up to 44 characters). The Remote Queue utility, FDRABRUT, will dynamically allocate this data set and add ARCHIVE RESTORE requests to it. One index level of the name (other than the first) must be 'ABRARCH' unless the DSNCK option has been set to NO. The default name is 'FDRABR.ABRARCH.DATA'; if you prefer to use a different name, change it here. [See Sections 51.07](#) and [90.45](#) for an explanation of the use of this data set by ABR.
- BKUPDMPQ** This is the fully-qualified name of the default BACKUP DUMP Remote Queue data set (up to 44 characters). The Remote Queue utility, FDRABRUT, will dynamically allocate this data set and add BACKUP DUMP requests to it, if the DISKUPDATE option is set to NO (if DISKUPDATE is YES, this data set is not used). One index level of the name (other than the first) must be 'ABRBKDQ' unless the DSNCK option has been set to NO. The default name is 'FDRABR.ABRBKDQ.DATA'; if you prefer to use a different name, change it here. [See Sections 50.04](#) and [90.45](#) for an explanation of the use of this data set by ABR.
- BKUPRSTQ** This is the fully-qualified name of the default BACKUP RESTORE Remote Queue data set (up to 44 characters). The Remote Queue utility, FDRABRUT, will dynamically allocate this data set and add BACKUP RESTORE requests to it. One index level of the name (other than the first) must be 'ABRREST' unless the DSNCK option has been set to NO. The default name is 'FDRABR.ABRREST.DATA'; if you prefer to use a different name, change it here. [See Sections 50.07](#) and [90.45](#) for an explanation of the use of this data set by ABR.
- NOTE:** If you plan to use any or all of the 4 remote queue datasets, you must pre-allocate and initialize them. [See Section 90.45](#).

90.17 ARCHIVE UTILITY DEFAULTS (PANEL A.I.4.6)

These options are defaults used by program FDRARCH, the ARCHIVE Control File Maintenance Utility (described in [Section 51.40](#)) and related ABR utilities. Most of these options can be overridden at execution time.

PANEL A.I.4.6

FDRARCH
DEFAULTS

----- FDR INSTALLATION -- SET ABR ARCHIVE UTILITY DEFAULTS -----	
COMMAND ==>	
FDRARCH UTILITY DEFAULTS	
BLKF	DEFAULT NUMBER OF BLOCKS PER TRACK IN THE ARCHIVE FILE 2
RECS	DEFAULT NUMBER OF DATA SET ENTRIES IN THE ARCHIVE FILE 2000
REORG%	PERCENTAGE OF FREE SPACE IN ARCHIVE FILE TO ISSUE WARNING 10
ARCDD	DDNAME USED BY FDRARCH TO REFERENCE ARCHIVE FILE ARCHIVE
DYNARC	DYNAMIC ALLOCATION OF ARCHIVE FILE BY FDRARCH ALLOWED NO
ERASE	ARCHIVE FILE MAY BE RE-FORMATTED (WRITTEN OVER) NO
ARCRESV	PROTECT ARCHIVE FILE WITH RESERVE DURING UPDATE FUNCTION YES
FDRARCH ARCHIVE FILE REORGANIZATION OPTIONS	
RECALL	UNCATALOG AUTO-RECALL DATA SETS THAT WERE DROPPED YES
UNCAT	UNCATALOG ARCHIVE TAPE FILES THAT ARE NO LONGER NEEDED YES
MSG	PRINT UNCATALOGED TAPE FILES THAT ARE NO LONGER NEEDED YES

BLKF This is the number of blocks per track (1 through 5) used when formatting the ARCHIVE Control File.

%REORG Whenever the percentage of free space remaining in the ARCHIVE Control File drops below this value (from 5 to 75), a PRINT ARCHIVE report (program FDRABRP, [section 53.03](#)) will issue a warning recommending reorganization. The default is 10.

ARCDD This is the DD name used by FDRARCH, FDRABRP, FDREPORT, and FDRTSEL, to refer to the ARCHIVE Control File. The default is 'ARCHIVE'. Note that ABR (program FDRABR) does **not** use this option; it always uses DDname ARCHIVE.

DYNARC If set to YES, FDRARCH and FDRTCOPY will dynamically allocate the default ARCHIVE Control File named in the ABR option table (Panel A.I.4.5) if the DD statement specified by "ARCDD" above is not present in the JCL. The default is NO. Note that this does not affect the DYNARC option of ABR (program FDRABR); the DYNARC option must be specified on ABR DUMP and RESTORE statements.

ERASE If set to YES, FDRARCH is permitted to re-format (overwrite) an existing ARCHIVE Control File. The default is NO.

ARCRESV If set to YES, FDRARCH will protect the ARCHIVE Control File, and application backup control files, with a RESERVE during any update. The default is YES, and Innovation recommends that you do not change it (damage to the Archive Control File may otherwise result unless you can be certain no other ABR process will be updating it concurrently). However, if you have shared DASD, Innovation also recommends that you use a cross-system ENQ facility to convert these RESERVEs to global ENQs; details are in the ENQ member of the FDR Installation Control Library (ICL).

RECALL If set to YES, an FDRARCH REORG function will uncatalog the auto-recall catalog entry for any data set being deleted from the ARCHIVE Control file by the REORG. If NO, auto-recall catalog entries will be left even though the data set has been deleted from the Control File. The default is YES.

UNCAT If set to YES, an FDRARCH REORG function will uncatalog any cataloged ARCHIVE backup data sets that are being deleted from the ARCHIVE Control file by the REORG (if the TMS option is enabled on panel A.I.4.4 and not overridden by a REORG option, REORG will only uncatalog backups whose expiration is recorded as 99000 (tape management catalog control)). If NO, cataloged ARCHIVE backups may be left even though there are no longer referred to by the Control File. The default is YES.

90.17 CONTINUED

MSG If set to YES, an FDRARCH REORG function will print a message about any cataloged ARCHIVE backup data sets being uncataloged because of the UNCAT option above. If NO, REORG will not display the backups that it has uncataloged. The default is YES.

90.18 ABR DISK PROCESSING OPTIONS (PANEL A.I.4.7)

These options are defaults used by FDRABRM, the ABR VTOC Maintenance Utility (described in [Section 50.40](#)) and the ABR ISPF panel A.I.8 (SET ABR VOLUME PROCESSING OPTIONS). Most of these can be overridden at execution time. It also includes options relating to ABR OLDBACKUP processing.

**PANEL A.I.4.7
ABR VOLUME
PROCESSING
OPTIONS**

```

----- FDR INSTALLATION -- SET ABR DISK VOLUME PROCESSING OPTIONS ---
COMMAND ===>

      OLD BACKUP TABLE OPTIONS
OLDBENT MAXIMUM NUMBER OF DATA SET OLD BACKUPS TABLED IN DSCB ..... 13
OLDBOFF DSCB RESERVED FIELD OFFSET TO USE FOR OLD BACKUP TABLE ..... 62

      FDRABRM UTILITY DEFAULTS
CYCLE   MAXIMUM NUMBER OF AUTO-CYCLES CREATED PER GENERATION ..... 10
GEN     NUMBER OF FULL VOLUME BACKUPS TO MAINTAIN CATALOGED ..... 4

RETPD   RETENTION PERIOD FOR ABR FULL VOLUME BACKUP TAPES (COPY1) ..... 60
RETP2   RETENTION PERIOD FOR ABR FULL VOLUME BACKUP TAPES (COPY2) ..... DITTO

FORCE   INITIALIZE DSCBS EVEN IF RESERVED BYTES ARE NOT ZERO ..... NO

CSLREF  IF LAST REFERENCE DATE IS ZERO, RESET IT TO CURRENT DATE ..... NO

LIST    PRODUCE VTOC REPORT OF DSCBS CHANGED BY VTOC MAINTENANCE ..... NO

```

**OLDBENT
OLDBOFF**

These are the OLD BACKUP options. The ABR Incremental Backup system always records the most current backup cycle on which a given data set was backed up; this value is stored in a field in the Format 1 DSCB reserved for ABR use. In addition, ABR provides the option of recording up to 13 previous backups in another DSCB field. If this option is used, then a user who wants to restore a data set from a backup that is not the most recent can easily display and select one of the older backups. If this option is not used, the user would have to know and specify the ABR GENERATION and CYCLE of the older backup.

The OLDBENT and OLDBOFF options specify a field within each Format 1 DSCB where the old backup information is to be stored. One byte is required for each old backup recorded. OLDBENT (the number of entries) is the length of the field in bytes; OLDBOFF is the displacement of the beginning of the field, in bytes, from the beginning of the Format 1 DSCB (relative to the beginning of the data set name). This must be a DSCB field which is not used for other purposes and is not stored into by other than ABR.

WARNING: at one time there were several such DSCB fields available for ABR use, but now the only safe field is DS1SYSCD (the "system code" field; this is initialized to a constant by allocation but is not used for any purpose). DS1SYSCD is 13 bytes long (OLDBENT=13) and is at displacement 62 (OLDBOFF=62). Innovation does not recommend using any other DSCB field. If your installation is using another field (any other values) for OLD BACKUP, contact Innovation for assistance in moving to DS1SYSCD.

It is also possible that you have another software product that is using all or part of DS1SYSCD (particularly other DASD management products). If so, contact Innovation for assistance in setting OLD BACKUP options correctly to avoid loss of information.

The OLDBACKUP option is activated on a volume by volume basis by the ABR ISPF panel A.I.8 (SET ABR VOLUME PROCESSING OPTIONS), or by the ABRINIT or MAINT commands of program FDRABRM ([Section 50.40](#)). These functions set an option flag enabling OLDBACKUP in the ABR MODEL DSCB on each volume. Once the option is enabled on a particular volume, ABR backups will begin recording data in the field specified by OLDBENT and OLDBOFF.

CYCLE

Specifies the default value for CYCLE (auto-cycles) for the ABRINIT and REMODEL commands of FDRABRM and with ISPF panel A.I.8. Default is 10. Note that CYCLE is used only by ABR backups specifying TYPE=AUTO; it does not limit the number of cycles that can be created in a given generation.

90.18 CONTINUED

- GEN** Specifies the default value for GEN (maximum generations to keep in the ABR catalog) for the ABRINIT and REMODEL commands of FDRABRM and with ISPF panel A.I.8. Default is 4.
- RETPD** Specifies the default values for RETPD (default retention of COPY1 full-volume backups) to be used with the ABRINIT and REMODEL commands of FDRABRM and with ISPF panel A.I.8. Default is 60.
- RETP2** Specifies the default values for RETP2 (default retention of COPY2 full-volume backups) to be used with the ABRINIT and REMODEL commands of FDRABRM and with ISPF panel A.I.8. A value of "DITTO" means to use the value of RETPD (COPY1 default) for COPY2 as well. Default is DITTO.
- FORCE** If set to YES, allows the ABRINIT command of FDRABRM and ISPF panel A.I.8 to initialize volumes for ABR processing even if some Format 1 DSCBs have non-zero data in the reserved bytes 103-104 used by ABR. Default is NO. **Innovation does not recommend setting this option since it can be specified when needed at execution time.**
- CSLREF** If set to YES, causes the ABRINIT and MAINT commands of FDRABRM and ISPF panel A.I.8 to set the last reference date on any data set with a zero LRDATE to today's date. Default is NO.
- LIST** If set to YES, causes all functions of FDRABRM and ISPF panel A.I.8 to list the names of all datasets changed by the function. Default is NO.

90.19 MORE ABR GENERAL OPTIONS (PANEL A.I.4.8)

Here are a few more options used by ABR and FDRABRM.

**PANEL A.I.4.8
MORE ABR
GENERAL
OPTIONS**

```
----- FDR INSTALLATION -- SET FDR GLOBAL OPTIONS FOR ABR -----
COMMAND ==>

ARCHI      DEFAULT ARCHIVE HIGH THRESHOLD (PERCENTAGE OF USED SPACE). 80

ARCLOW     DEFAULT ARCHIVE LOW  THRESHOLD (PERCENTAGE OF USED SPACE). 50

MAXFILES   MAXNUMBER OF FILES TO CREATE IN A MULTI-FILE BACKUP TAPE. 255
```

ARCHI This is the default value (0 to 99) for the high-allocation threshold used by ARCHIVE and SUPERSCRATCH if THRESHOLD=HIGH is specified on the DUMP statement ([see Section 51](#)). This threshold is stored in the ABR model, so this is the default used by FDRABRM and panel A.I.8 (although it can be overridden). This default is also used by ABR for volumes on which a threshold has never been set. Default is 80.

ARCLOW This is the default value (0 to 99) for the low-allocation threshold used by ARCHIVE and SUPERSCRATCH if THRESHOLD=LOW is specified on the DUMP statement ([see Section 51](#)). This threshold is stored in the ABR model, so this is the default used by FDRABRM and panel A.I.8 (although it can be overridden). This default is also used by ABR for volumes on which a threshold has never been set. Default is 50.

MAXFILES This is the default for the MAXFILE= operand in ABR backups. Unless overridden, it specifies the maximum number of files that ABR will create on a multi-file backup tape, from 1 to 4095. If you will be outputting to high-capacity tapes, such as IBM Magstar 3590 or STK Redwood, you may want to increase this value. Default is 255.

90.20 REPORT OPTIONS (PANEL A.I.4.9)

This panel (and the following panel A.I.4.10) contain defaults used by the FDRABRP and FDREPORT report programs, described in [Sections 53](#) and [54](#). These option values can be overridden at execution time. FDREPORT also has some defaults used only by FDREPORT; see [Section 54](#) for details.

**PANEL A.I.4.9
REPORT
OPTIONS**

----- FDR INSTALLATION -- SET FDR GLOBAL REPORT OPTIONS ----- COMMAND ==>	
ARCHIVE REPORT	
SDAYS	NUMBER OF DAYS BACK TO CALCULATE STOP DATE 0
XDAYS	NUMBER OF DAYS AHEAD TO CALCULATE FUTURE EXPIRATION DATE 10
ARCSUM	PRINT ARCHIVE FILE SUMMARY YES
VTOC REPORT	
DETAIL	PRINT DATA SET INFORMATION IN VTOC LIST FUNCTIONS YES
SUM	PRINT SUMMARY IN VTOC LIST FUNCTIONS YES
SORT OPTIONS	
SORTCORE	STORAGE REQUIRED FOR EXTERNAL SORT (BYTES) 100000
SORTMSG	MESSAGE OPTION FOR EXTERNAL SORT PROGRAM CC
SORTPFX	DDNAME PREFIX FOR EXTERNAL SORT PROGRAM SORT
LASTABR	DAYS AFTER ABR LAST PROCESSED THE VOLUME TO ISSUE WARNING 7
OLDBACKUP	OLD BACKUP TO REPORT.....(CURRENT)

SDAYS Specifies a value in days (0 to 9999) used to calculate a "stop date" (run date minus SDAYS) when processing the ARCHIVE Control File. Only data sets that were ARCHIVED on or after that will be printed. This option can reduce the execution time of the print program if many data sets have been archived. The default (0) disables the option (all data sets are printed).

XDAYS Specifies a value in days (0 to 9999) used to calculate a future expiration date when processing the ARCHIVE Control File. This date is used when calculating the number of expired entries printed in the summary report. This value is also used if the user selects expired entries by coding EXPIRE operand. The default is 10.

ARCSUM If set to YES, the archive control file summary report will be produced during a PRINT ARCHIVE function. If set to NO, the summary is suppressed. The default is YES.

DETAIL If set to YES, data set information will be printed by a VTOC list function. If set to NO, the data set information is suppressed. The default is YES.

SUM If set to YES, if the summary will be produced after each VTOC list. If set to NO, the summary is suppressed. The default is YES.

SORTCORE Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be from 10000 to 8000000 inclusive. The default is 100000.

SORTMSG Specifies the message option to be used by the program 'SORT' if external sorting is required.

AC -- all messages to the console	AP -- all messages to the printer (SYSOUT)
CC -- critical messages to the console	CP -- critical messages to the printer
NO -- no messages to be produced	PC -- critical messages to both console and printer

The default is CC.

90.20 CONTINUED

SORTPFX Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used. The default is SORT.

LASTABR When reporting on volume status, if the report program detects that a volume has not been processed by ABR backups in the number of days specified (from 1 to 90), a warning message is issued. This helps to identify volumes which may have accidentally been bypassed. The default is 7 (one week).

OLDBACKUP Specifies which backups will be printed when the cross reference (XREF) option of 'PRINT BACKUP' or 'PRINT SCRATCH' function of FDRABRP is invoked. A value of "ALL" prints all backups associated with the data set. "CUR" or "CURRENT" prints the current backup only. Numeric values of 0 to 13 inclusive print specific relative backup numbers, where 0 is the current backup and 13 is the oldest possible backup tape (multiple numeric values can be specified, separated by commas, e.g., "0,3,7").

90.21 MORE REPORT OPTIONS (PANEL A.I.4.10)

These are additional defaults for FDREPORT and FDRABRP.

Recommendation: Most users will have no need to change any of these defaults, except possibly MAXONLINE.

**PANEL A.I.4.10
MORE REPORT
OPTIONS**

```

----- FDR INSTALLATION -- SET FDR GLOBAL REPORT OPTIONS -----
COMMAND ==>

ICFSOURCE  SOURCE OF ICF VSAM DATA AND INDEX COMPONENT INFORMATION... DEFAULT
            ( DEFAULT  VVDS  LOCATE )

MAXICF      MAXIMUM NUMBER OF ICF CLUSTERS PROCESSED AT ONE TIME ..... 400

GBLDSN      GENERIC NAME FOR SELECTION OF ALL DATA SETS ..... &ALLDSN

GBLVOL      GENERIC NAME FOR SELECTION OF ALL VOLUMES ..... &ALLV

DYNAL       DYNAMIC ALLOCATION OF DISK DEVICES ALLOWED ..... YES
MAXONLINE   MAXIMUM NUMBER OF VOLUMES PROCESSED PER PROGRAM EXECUTION ..... 256

MAXTCVX     MAXIMUM NUMBER OF BACKUP TAPES TABLED FOR ALL DISK VOLUMES ..... 1000
MAXTSVX     MAXIMUM NUMBER OF BACKUP TAPES TABLED PER DISK VOLUME ..... 100

XREFERROR   LOW SEVERITY WARNING MESSAGES ARE PRINTED IN BACKUP REPORT ..... YES

```

- ICFSOURCE** Specifies the source of information for reporting on ICF VSAM data and index components.
DEFAULT - if the program runs APF authorized, read the VVDS directly (see VVDS option value). Otherwise, issue the LOCATE SVC to retrieve the necessary information.
LOCATE - always issue the LOCATE SVC to retrieve the necessary information.
VVDS - read the 'SYS1.VVDS.Vvolser' cluster directly and extract the necessary information directly from the VVR records. The program must run APF authorized.
- MAXICF** Specifies the maximum number of ICF clusters whose information retained in storage during a single program execution. The number may be from 100 to 10000 inclusive. The default is 400.
- GBLDSN** Specifies a character string from 1 to 7 characters in length that can be used after DSN= to specify all data sets. An ampersand(&) must prefix the specified string. The default is &ALLDSN (e.g., use DSN=&ALLDSN to select all data sets on a volume).
- GBLVOL** Specifies a character string from 1 to 5 characters in length that can be used after VOL= to specify all online disk volumes. An ampersand(&) must prefix the specified string. The default is &ALLV (e.g., use VOL=&ALLV to select all online volumes).
- DYNAL** If set to YES, dynamic allocation of disk devices is permitted. If set to NO, DISKx DD statements are required. The default is YES.
- MAXONLINE** Specifies the maximum number of disk volumes that may be allocated via DD statements and/or by dynamic allocation during any single program execution. The number may be from 5 to 9000 inclusive (Note: Program FDRABR also uses MAXONLINE but will use 256 if the MAXONLINE value is less than 256). The default is 256.
- MAXTCVX** Specifies the maximum number of tapes that will be tabled for all disk volumes during the execution of a combined DSNAMES/BACKUP tape cross reference (XREF) report. The number may be from 100 to 32000 inclusive. The default is 1000.
- MAXTSVX** Specifies the maximum number of tapes that will be tabled for a given disk volume during the execution of a volume oriented DSNAMES/BACKUP tape cross reference (XREF) report under FDRABRP, or a report that includes backup information under FDREPORT. The number may be from 10 to 300 inclusive. The default is 100.
- XREFERROR** If set to YES, permits the printing of warning messages when the cross reference (XREF) of 'PRINT SCRATCH' or 'PRINT BACKUP' is selected. If set to NO, warning messages be printed only if no other backup tape messages are to be printed. The default is YES.

90.22 OPERATING SYSTEM EXIT OPTIONS (PANEL A.I.4.11)

The FDR system includes several Operating System exits for FDRABR and FDRREORG. Although these exits are not required for ABR and FDRREORG to function, they provide very important additional functions. Additional information on the exits is found in [Section 90.60](#). This panel takes you to other panels (shown below) where you can indicate which exits to install, and set options for the various exits.

PANEL A.I.4.11
OPERATING
SYSTEM EXIT
OPTIONS

```
----- FDR INSTALLATION -- SET OPERATING SYSTEM EXIT OPTIONS ----
OPTION  ==>

  1  - DYNAMIC INSTALLATION EXITS

  2  - ABR CATALOG LOCATE EXIT OPTIONS

  3  - MORE ABR CATALOG LOCATE EXIT OPTIONS

  4  - ABR DATA SET NOT FOUND EXIT OPTIONS

  5  - ABR DADSM PRE-PROCESSING EXIT OPTIONS
```

PANEL A.I.4.11.1
DYNAMIC
INSTALLATION

This panel sets options for the dynamic installation of the ABR exits. The options will control the operation of FDRSTART, the FDR utility which controls installation of the exits ([see Section 90.60](#)).

```
----- FDR INSTALLATION -- SET DYNAMIC INSTALLATION EXITS -----
COMMAND ==>

ABRLOC   DYNAMICALLY INSTALL THE ABR CATALOG LOCATE EXIT ..... NO
ABRDSNF  DYNAMICALLY INSTALL THE ABR DATASET NOT FOUND EXIT ..... NO
ABRPRE   DYNAMICALLY INSTALL THE ABR DADSM PRE-EXIT ..... NO

IEBCOPY  DYNAMICALLY INSTALL FDRREORG TO PROCESS IEBCOPY COMPRESS ..... NO

CONFINST CONFIRM INSTALLATION OF THE FDR SYSTEM EXITS AT IPL TIME ..... NO
```

ABRLOC If set to YES, the ABR CATALOG LOCATE exit will be dynamically installed. The LOCATE exit is **required** to automatically recall (auto-recall) ARCHIVED data sets. The default is NO.

ABRDSNF If set to YES, the ABR Data Set Not Found (DSNF) exit will be dynamically installed. The DSNF exit is required for auto-recall of ARCHIVED data sets not referenced through the catalog. The default is NO.

ABRPRE If set to YES, the ABR DADSM Preprocessing exit will be dynamically installed. The DADSM Preprocessing exit performs several functions for ABR, including the recording of scratched and renamed data sets. **Its use is strongly recommended for all ABR installations.** The default is NO.

IEBCOPY If set to YES, the FDRREORG IEBCOPY intercept will be dynamically installed. The IEBCOPY intercept routes IEBCOPY compress requests to FDRREORG. This option should be used only if the FDRREORG product is installed. The default is NO.

CONFINST If set to YES, FDRSTART will ask the operator to confirm the dynamic installation of the exits via a console message (WTOR). The default is NO (no confirmation is requested).

90.22 CONTINUED

**PANEL
A.I.4.11.2
ABR AUTO-
RECALL
OPTIONS**

The following options affect only the CATALOG LOCATE exit, except for LXBYPSEC which also applies to the Data Set Not Found (DSNF) exit.

----- FDR INSTALLATION -- SET ABR GLOBAL AUTO-RECALL OPTIONS -----	
COMMAND ==>	
LXFOREST	USER IS ASKED TO CONFIRM RESTORE..... YES
LXDFREST	RESTORE TYPE(S) FOR TSO AUTO-RECALL..... FG,BG,RQ (NO NONE FG FOREGROUND BG BACKGROUND RQ REMOTE QUEUE)
LXFGSYNBG	CONVERT TSO FG RESTORE TO SYNCHRONOUS BACKGROUND RESTORE..... NO
LXFGERR	RETRY AUTO-RECALL REQUEST IF FG FAILS (NO, BG OR RQ)..... NO
LXNOMSG	SUPPRESS AUTO-RECALL MESSAGES IF TSO RECALL IS DISABLED..... NO
LXNEWVOL	USER MAY SPECIFY A NEW VOLUME SERIAL WHERE TO RESTORE..... YES
LXREISSU	PROMPT TO CONFIRM THE SPECIFIED NEW VOLUME SERIAL..... NO
LXSYNPROC	CATALOGED PROCEDURE NAME USED FOR BACKGROUND RESTORES..... SYNRECAL
AUTO-RECALL SECURITY OPTIONS	
LXCHKSEC	CHECK SECURITY BEFORE THE AUTO-RECALL RESTORE OPERATION..... NO
LXBYPSEC	BYPASS SECURITY DURING AUTO-RECALL RESTORE OPERATION..... NO
LXEXIT	RECALL SECURITY EXIT ACTIVE..... NOEXIT NAME..... NONE

LXFOREST

If set to YES, when the CATALOG LOCATE exit finds that a TSO user has referenced a data set that is ARCHIVEd for recall, it will ask the user to confirm (via message FDRW71) that the data set should be recalled. If NO, the recall proceeds automatically without asking for confirmation. This option has no effect if options LXDFREST and LXUNCAT are both set to NO. The default is YES.

LXDFREST

Specifies the CATALOG LOCATE exit restore options that are available to do automatic recalls for TSO users. The values that may be specified are NO or any combination of FG, BG, and RQ. If more than one restore type is specified, the TSO user will be prompted to choose which restore type to use. If only one restore type is specified, that restore type will be used automatically, and the TSO user will not be prompted. The restore types are described in [section 51](#).

NO -- specifies that automatic recall is not available for TSO users.

FG -- foreground recalls (immediately under TSO or as a synchronous external task) are available.

BG -- background recalls (asynchronous external task) are available

RQ -- remote queue recalls (added to remote queue for later processing) are available.

The default is FG,BG,RQ (all restore types).

Recommendation: remove RQ unless you have implemented remote queue processing (see [Section 90.45](#)). If you ARCHIVE only to tape, remove FG. But if you create one archive copy on disk, you may want to leave FG enabled so that TSO users can get quick synchronous recalls; in this case you will probably want to set LXFGERR to BG or RQ (see below).

LXFGSYNBG

If set to YES, any requests for TSO foreground recalls via the CATALOG LOCATE exit are actually processed as synchronous background recalls. The recall is done by a separate started task (similar to background recalls) but the TSO user will wait until the recall is complete. This may be useful when TSO users do not have tape mount privileges but it may mean that TSO users will have their sessions locked up longer while the recall takes place. The default is NO.

90.22 CONTINUED

LXFGERR If set to BG or RQ, the CATALOG LOCATE exit will retry a recall in the background (BG) or via the remote queue (RQ), if it attempted in the TSO foreground but fails. This option has no effect unless option LXDFREST includes FG as an available restore type. The default is NO (failed foreground recalls are not retried).

LXFGERR is an important option in the environment where data sets are ARCHIVED to disk with a short retention and to tape with a longer retention, and the TSO users do not have the authority to mount tapes (MOUNT). As long as an ARCHIVED data set is available on disk, it can be quickly be recalled in the foreground. After the disk copy expires, ABR will automatically try to restore from the tape copy, but the dynamic allocation will fail because the user does not have MOUNT privileges. If LXFGERR is set to BG or RQ, the recall will automatically be sent to the background or the remote queue, where the tape can be mounted.

Note that you can specify BG or RQ even if BG or RQ is not specified as a valid restore type under LXDFREST. For example, if you specify FG as the only value for LXDFREST, and you specify BG as the value for LXFGERR, then all recalls will initially be tried in the foreground, and those that fail will automatically be retried via START commands.

LXNOMSG If set to YES, the CATALOG LOCATE exit will suppress messages FDRW70 and FDRW79, when the exit finds that the user has referenced a data set that is ARCHIVED for recall, but the installation does not allow automatic recalls for TSO users (LXDFREST and LXUNCAT both set to NO). The default is NO (issue the messages for a auto-recall data set if TSO auto-recall is disabled).

LXNEWVOL If set to YES, the CATALOG LOCATE exit will permit the TSO user to designate an output volume to which to restore an ARCHIVED data set with messages FDRW76 (original volume serial) and FDRW77 (specify new output volume serial).

The default is YES. However, in many installations users do not get involved in choosing specific output volumes, and the selection of output volumes may be automated by the ABR Restore Allocate List, SMS, or other products, so you will probably want to change this to NO.

The topic "Output Volume for Recall" at the end of this section discusses the interaction among the various options that may affect the output volume choice.

Recommendation: change LXNEWVOL to NO.

LXREISSU If set to YES, once the CATALOG LOCATE exit has asked the user for a new output volume serial (see LXNEWVOL), it will re-display the designated volume and give the user another chance to change it; this will continue until the user accepts the last entry by pressing ENTER. This option has no effect if either option LXDFREST or option LXNEWVOL is set to NO. The default is NO (user gets only once chance to designate the output volume).

LXSYNPROC Specifies the name of the cataloged procedure that will be executed as a started task by the CATALOG LOCATE exit to perform external restores. External recalls are used under TSO for asynchronous background recalls (recall type BG) and synchronous background recalls (see option LXFGSYNBG above). They are also used for recalls for batch jobs in special circumstances when a recall in the batch address space is not possible. The default is SYNRECAL.

When installing the CATALOG LOCATE exit, you must copy the SYNRECAL cataloged procedure from the Installation Control Library (ICL) to a cataloged procedure library that is available to JES for START commands. You may change the name of this cataloged procedure from 'SYNRECAL' to a name of your choice; if so, specify the new name as the value of the LXSYNPROC option. Further details are in [Sections 90.61](#) and [90.62](#).

90.22 CONTINUED

- LXCHKSEC** If set to YES, the CATALOG LOCATE exit itself will perform data set security checking for READ authority before recalling an ARCHIVED data set. This is in addition to security checking that may be done by FDRABR during the restore or by FDRABRUT during the remote queue update. "Security Considerations" in [Section 51](#) gives further details. This option has no effect if global option ALLCALL is set to NO. The default is NO (no security checking in the LOCATE exit).
- Recommendation:** set LXCHKSEC to YES if you have a security system.
- LXBYPSEC** If set to YES, the CATALOG LOCATE exit and the Data Set Not Found (DSNF) exit will cause security checking to be bypassed during the ABR restore. This option is separate from the LXCHKSEC and LXEXIT options, which control security checking that may be done by the CATALOG LOCATE exit itself before FDRABR or FDRABRUT is invoked. The bypass applies both to security checking that FDRABR and FDRABRUT perform explicitly ([see Section 90.12](#)), and to security checking performed by Operating System routines that FDRABR and FDRABRUT invoke. With LXBYPSEC=NO, users may require ALTER authority to allocate and recall an ARCHIVED data set; LXBYPSEC=YES allows users with READ or no authority to recall the data set, although normal security will be in place again when they attempt to open it. LXBYPSEC=YES and LXCHKSEC=YES will require at least READ authority to recall a data set. "Security Considerations" in [Section 51](#) gives further details.
- Recommendation:** set LXBYPSEC to YES if you have a security system.
- LXEXIT** If set to YES, the CATALOG LOCATE exit will invoke a special security exit for each data set to be recalled (EXIT NAME specifies the name of the load module containing the exit). The exit may also be used to specify a new output volume for the recalled data set. The topic "Output Volume for Recall" at the end of this section discusses the interaction among the various options that may affect the output volume choice. The default is NO. There is no default for the name of the exit; if you activate the exit you must specify its name. If you wish to write a security exit, please contact Innovation for details.

90.22 CONTINUED

The following options affect only the CATALOG LOCATE exit.

**PANEL
A.I.4.11.3
MORE ABR
AUTO-RECALL
OPTIONS**

```

----- FDR INSTALLATION -- MORE ABR GLOBAL AUTO-RECALL OPTIONS -----
COMMAND ===>

LXNCDENY    FAIL LOCATE WITH 'NOT CATALOGED' IF TSO USER DENIES RECALL YES

LXCONUSE    USE CONSTANT NEW VOLUME FOR ALL DATA SET RECALLS..... NO
LXCONVOL    CONSTANT NEW VOLUME FOR ALL DATA SET RECALLS..... NEWVOL

LXALTMSG    ISSUE ALTERNATE FDRW71 (CONFIRM RESTORE) MESSAGE FORMAT... NO

LXUNCAT     ASK USER WHETHER TO UNCATALOG INSTEAD OF RECALLING DATASET NO

LXMAXSTC    MAXIMUM NUMBER OF RECALL STARTED TASKS ACTIVE AT ONE TIME. NONE
LXMAXREC    LIMIT OF TOTAL RECALLS (INCLUDING STC) ACTIVE AT ONE TIME. NONE

LXDIRVOL    DIRECTED VOLUME SERIAL NUMBER TO BE RETURNED TO JES3..... NONE
LXDIRTYP    DIRECTED DEVICE TYPE TO BE RETURNED TO JES3..... NONE

LXSPFMIG    DISPLAY VOLSER MIGRAT FOR MULTIVOL DATASETS UNDER ISPF 3.4 NO

```

LXNCDENY If set to YES, the CATALOG LOCATE exit will give a 'not cataloged' return code to LOCATE when a TSO user references a data set that has been ARCHIVED for recall, and the TSO user specifies that the data set should not be recalled (by the reply to message FDRW71). This option has no effect if options LXDFREST and LXUNCAT are both set to NO. If set to NO, when the user bypasses the recall, a normal (zero) return code will be passed to LOCATE, along with the volume serial currently in the catalog (which may be MIGRAT); since the data set has not been recalled, TSO will probably try to open the data set and get an error saying it could not be found, or the volume could not be mounted. The default is YES.

LXCONUSE If set to YES, the CATALOG LOCATE exit will pass the constant new volume serial specified by the LXCONVOL option as the target output volume for automatic recall of every data set. If set to NO, the volume to which the data set is currently cataloged will be passed as the target output volume. The default is NO.

The usual use of the LXCONUSE/LXCONVOL options is to designate a non-existent volume in order to force ABR to select a volume using the ABR Restore Allocate List ([see Section 90.32](#)). If LXCONUSE is NO, ABR will attempt to restore to the original volume before trying volumes in the Restore Allocate List. The topic "Output Volume for Recall" at the end of this section discusses the interaction among the various options that may affect the output volume choice. **Specify YES to force use of the Restore Allocation List for all auto-recalls.**

LXCONVOL Specifies the constant new volume serial that the CATALOG LOCATE exit will specify as the output volume for automatic recall, if the LXCONUSE option is set to YES.

LXALTMSG If set to YES, the CATALOG LOCATE exit will use an alternate format for message FDRW71, when asking a TSO user whether an ARCHIVED data set should be recalled; this alternate format requires a positive action (keying in 'YES') to cause an ARCHIVED data set to be recalled. If set to NO, the standard format for FDRW71 is used, where pressing "Enter" recalls the data set and a response of "NO" is required to bypass the recall. This option has no effect if options LXDFREST and LXUNCAT are both set to NO, or if option LXFOREST is set to NO. The default is NO.

90.22 CONTINUED

LXUNCAT If set to YES, the CATALOG LOCATE exit will offer a TSO user the choice of uncataloging an archived data set instead of recalling; this allows the user to create a new data set if the data in the archived version is not required. If option LXFOREST is set to NO, then option LXUNCAT will be ignored, and the TSO user will not be offered the choice of uncataloging the data set. The default is NO (do not offer option of uncataloging).

LXMAXSTC Specifies the maximum number of recall started tasks that the CATALOG LOCATE exit will allow to be running in the system at one time. The number may be from 1 to 255. If LXMAXSTC and LXMAXREC are both specified, then LXMAXSTC must be less than or equal to LXMAXREC. If set to NONE, the exit will not limit the number of recall tasks.

A Started task is used for recall for background (BG) recalls, for foreground recalls if LXFGSYN BG is set to YES, and for recalls for batch jobs for some special situations; the proc name specified by LXSYNPROC is used for these tasks. If the limit is reached, then:

-- If a batch job is requesting the recall, the CATALOG LOCATE exit will issue a message to the console operator, and will go into a wait for one minute at a time, until the number of started tasks for recall falls below the limit. The console operator can cancel the job while it is waiting.

-- If a TSO user is requesting the recall, the CATALOG LOCATE exit will issue a message to the TSO terminal, asking the user whether he wants to wait, or to terminate the current operation and perhaps try again later. If the TSO user chooses to wait, then the CATALOG LOCATE exit will go into a wait for one minute at a time, until the number of started tasks for recall falls below the limit. The TSO user can use the Attention (PA1) key to terminate the wait if he changes his mind.

The STC count is incremented by the CATALOG LOCATE exit when a task is started, and decremented by the task when it is completed. If the count becomes incorrect (which can occur when recall tasks are cancelled or fail to start), the LOCATE exit will automatically correct the count shortly after the maximum is reached.

Note: you may need to increase the maximum permissible number of address spaces in your system to accommodate the recall started tasks (see the parameters MAXUSER= and RSVSTRT= in the IEASYSxx member of SYS1.PARMLIB).

LXMAXREC Specifies the maximum total number of recalls that the CATALOG LOCATE exit will allow to be running in the system at one time. This is the sum of recalls done by started tasks (which may be limited separately by the LXMAXSTC option) and recalls performed in the address space issuing the LOCATE. (Remote queue recalls do not count.) The number may be from 1 to 255. If LXMAXSTC and LXMAXREC are both specified, then LXMAXSTC must be less than or equal to LXMAXREC. If exceeded, the result is the same as described for LXMAXSTC above. If set to NONE, the exit will not limit the total number of recalls.

LXDIRVOL (JES3 only) To meet the requirements of JES3 job scheduling, this specifies a special DASD volume serial which will be returned to JES3 during setup for all data sets archived for auto-recall. This "directed" volume serial must be online and available to all systems in the complex so that JES3 will be able to schedule the job on any system (subject to other constraints). Once the job begins running and the data set is recalled, the actual volume serial to which it was recalled will be used.

LXDIRTYP (JES3 only) Specifies the device type of the "directed" volume serial in LXDIRVOL above, specified as "3380", "3390", etc. This must be the actual device type on which that volume is mounted.

LXSPFMIG Specifies that under ISPF option 3.4 (data set list utility), the volume serial of data sets archived for auto-recall will be displayed as "MIGRAT".

Note: for ISPF Version 4.1 and later, this option is required only for multi-volume data sets, and works only when the "TOTAL" display view is requested.

90.22 CONTINUED

**OUTPUT
VOLUME FOR
RECALL**

Whenever the CATALOG LOCATE exit initiates a restore for an ARCHIVED data set, using any restore type, the CATALOG LOCATE exit designates an output volume to which the data set should be restored. The default is that the designated volume is the volume to which the data set is cataloged (which is either the volume from which the data set was ARCHIVED, or 'MIGRAT' if the MIGRAT option was in effect when the data set was ARCHIVED). The following paragraphs explain the interaction among the various options that may change the designated volume. These options are LXCONUSE, LXCONVOL, LXCHKSEC, LXEXIT, LXNEWVOL, and LXREISSU.

If the LXCONUSE option is set to YES, the CATALOG LOCATE exit initially set the output volume to the value specified by the LXCONVOL option. If not, it is set to the volume to which the data set is cataloged (which is either the volume from which the data set was ARCHIVED, or 'MIGRAT' if the MIGRAT option was in effect when the data set was ARCHIVED).

Next, if the LXEXIT option is set to YES, the CATALOG LOCATE exit calls the user-written recall security exit named by the EXIT NAME option. The LXEXIT exit may designate a new output volume. If the LXCHKSEC is set to YES, the CATALOG LOCATE exit may call a user-written FDR Data Set Security exit which may also designate a new output volume.

If the recall is on behalf of a TSO user, and the LXNEWVOL option is set to YES, the CATALOG LOCATE exit issues messages FDRW76 and FDRW77, telling the TSO user the output volume that is currently selected, and allowing entry of a new volume. If a new volume was entered, and the LXREISSU option is set to YES, the FDRW76/77 messages are issued again, allowing the user to change the volume again (this continues until "Enter" is pressed to accept the last choice).

Note that it would be very unusual for you to have all of the above options set to YES. For example, if your installation uses an LXEXIT exit to designate the output volume or uses SMS or the ABR Restore Allocation List to control output volume selection, you would probably want to set the LXNEWVOL option to NO so that the TSO user could not specify an output volume.

After all this, the output volume chosen by the CATALOG LOCATE is not necessarily the volume to which the data set will actually be restored. If the designated volume is 'MIGRAT', ABR will substitute the original volume from which the data set was ARCHIVED, as recorded in the ARCHIVE control file. If the designated or original volume is not online, or has no room for the data set, ABR will use the Restore Allocate List ([Section 90.32](#)), if any, to select a different volume. LXCONUSE and LXCONVOL, in particular, are often used to designate a non-existent volume in order to force ABR to select a volume using the Restore Allocation List. Also, if SMS is active, SMS may assign the data set to a different volume ([Section 70](#)).

90.22 CONTINUED

**PANEL A.I.4.11.5
ABR DADSM
PREPROCESSING
EXIT OPTIONS**

These options affect the ABR DADSM Preprocessing Exit.

----- FDR INSTALLATION -- SET ABR DADSM PRE-EXIT OPTIONS -----	
COMMAND ==>	
PXNOGDG	NUMBER OF GDG GENERATIONS TO KEEP IN THE SCRATCH CATALOG... 4
PXERRLEN	ISSUE WTO MESSAGE IF SCRATCHED DSNAME EXCEEDS 42 CHARACTERS YES

PXNOGDG

Specifies the number of scratch generations of a generation data group (GDG) that are to be kept in the ABR Scratch Catalog. PXNOGDG may be specified as an absolute number (1 to 255), or alternatively as a relative number, when prefixed by a plus or minus sign, e.g., -5, +12, +0. The first time that a generation of a given GDG is scratched, a GDG base will be built in the ABR scratch catalog with a name of "#.gdgbase" where "#" is the scratch catalog prefix specified on panel A.I.4.5 and "gdgbase" is the name of the actual GDG base; it will be assigned a limit value assigned by PXNOGDG so the records of that many scratched generations will be retained in the scratch catalog.

When PXNOGDG is a relative number, the limit value is calculated by taking the limit currently associated with the real GDG base and adding or subtracting the PXNOGDG value. For example, if PXNOGDG is +3, and a generation is scratched from a GDG having a limit of 5, the "#.gdgbase" in the scratch catalog is built with a limit of 8, so the GDG may have 5 live generations and 8 scratched generations.

Recommendation: Consider setting PXNOGDG to a higher value.

PXERRLEN

If set to YES, the DADSM Preprocessing exit will issue the WTO message (FDRW90 Reason=O) when it is unable to create an entry in the ABR Scratch Catalog for a data set being scratched or renamed because the combined lengths of the ABR Scratch index (#.) and of the data set name exceed 44 characters, i.e., when the data set name is 43 or 44 characters in length. If set to NO, no message is generated for this condition. The default is YES.

Innovation recommends that you set the PXERRLEN option to NO to avoid cluttering up the console with messages if you have many long data set names.

Note: Normally the records of scratched data sets with names of 43 or 44 characters are not recorded by the exit, leaving you with no easy way of identifying the backups of those data sets. If you have many such data sets, you can request custom zap C-53.0397 from Innovation; this will record those data sets but will truncate the last 1 or 2 characters from the name. The resulting entry cannot be used for automatic restores (because the name is not accurate) but users can display the recorded name and submit manual restores using the backup information.

90.23 SET FDRREORG OPTIONS (PANEL A.I.4.12)

This panel takes you to additional panels (described below) where defaults for the FDRREORG product can be set. Although all of these defaults can be overridden at execution time, specifying appropriate defaults here will greatly simplify your FDRREORG jobstreams. FDRREORG options are stored in the FDR program library in a module called FDRREOOP. **If you are not licensed for FDRREORG you can skip these panels.**

**PANEL A.I.4.12
FDRREORG
OPTIONS**

```

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----
OPTION  ===>

  1  -  BACKUP RELATED DEFAULTS FOR REORG COMMAND

  2  -  CHECKPOINT RELATED DEFAULTS FOR REORG AND RECOVER COMMANDS

  3  -  LOG FILE RELATED DEFAULTS FOR REORG AND RECOVER COMMANDS

  4  -  MISC. DEFAULTS FOR REORG, SIMULATE, AND RECOVER COMMANDS

  5  -  MORE MISC. DEFAULTS FOR REORG, SIMULATE, AND RECOVER COMMANDS

  6  -  DEFAULT SELECTION PARAMETERS FOR SELECT COMMAND

```

**PANEL
A.I.4.12.1
BACKUP
DEFAULTS**

```

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----
COMMAND  ===>

      BACKUP RELATED DEFAULTS FOR REORG COMMAND
BACKUP   BACKUP TYPE (PERM, GDG, TEMP)..... TEMP
BACKUPALLOC BACKUP ALLOCATION METHOD (SMS, UNIT)..... UNIT

BACKUPDATACLASSBACKUP SMS DATA CLASS NAME.....
BACKUPMGMTCLASSBACKUP SMS MANAGEMENT CLASS NAME.....
BACKUPSTORCLASSBACKUP SMS STORAGE CLASS NAME.....

BACKUPUNIT BACKUP UNIT NAME..... SYSDA
BACKUPUNITS BACKUP UNITS..... 1
BACKUPRETPD BACKUP RETENTION PERIOD..... 0
MAXFILE     MAXFILES ON BACKUP TAPE..... 255

BACKUPINDEX BACKUP INDEX.... ++BACKUP
LASTAPEPREFIX LASTAPE PREFIX..... FDRREORG
BACKUPFAILURE HOW TO TREAT A BACKUP FAILURE (ERROR, WARNING)..... ERROR
IAMCOMPPERM DO PERMANENT BACKUP OF IAM FILES IN COMPRESSED FORMAT NO
IAMCOMPTMP DO TEMPORARY BACKUP OF IAM FILES IN COMPRESSED FORMAT YES

```

BACKUP

Defines the default disposition of the backup data sets after the target data set has been successfully reorganized.

TEMP -- Backup data sets on DASD will be deleted and uncataloged. Backup data sets on TAPE will be uncataloged.

PERM -- Backup data sets on DASD will not be deleted and uncataloged. Backup data sets on TAPE will not be uncataloged.

GDG -- Same as PERM except backup data sets will be a +1 generation data set. FDRREORG will dynamically define any GDG bases that do not already exist. The high level qualifier of the dynamically generated generation data groups must be aliased to an ICF catalog.

The default is TEMP.

**BACKUP-
ALLOC**

Defines the allocation method to use when allocating backup data sets.

SMS -- Allocate by SMS storage class.

UNIT -- Allocate by unit name.

The default is UNIT.

Note: If set to SMS, a valid SMS storage class name must also be specified for BACKUPSTORCLASS. If set to UNIT, a valid unit name must also be specified for BACKUPUNIT.

90.23 CONTINUED

BACKUPDATA CLASS	Defines the default SMS data class for backup data sets when BACKUPALLOC is SMS.
BACKUP MGMTCLASS	Defines the default SMS management class for backup data sets when BACKUPALLOC is SMS.
BACKUP STORCLASS	Defines a valid SMS storage class which will be used to allocate SMS managed backup data sets when BACKUPALLOC is SMS.
BACKUPUNIT	Defines a valid unit name containing tape or disk devices to be used for allocating backup data sets when BACKUPALLOC is UNIT. It may be any value valid for UNIT= in JCL, except that the unit name specified must not contain a mixture of disk and tape devices. It may contain a mixture of device types (i.e. 3380s and 3390s). The default is SYSDA.
BACKUPUNITS	For backup data sets on tape, this defines the number of tape units to be allocated for backup data sets for each task. For backup data sets on disk, this operand defines the minimum number of disk volumes to be allocated for each backup data set; FDRREORG will dynamically increase the number of disk volumes to ensure that sufficient space is available to backup the data set being reorganized. The default is 1.
BACKUPRETPD	Defines the default retention period (equivalent to the JCL parameter RETPD=) to be used for all backup data sets.
MAXFILE	Specifies the maximum number of backup files (1-4095) to create on a tape volume. The default is 255.
BACKUPINDEX	Defines the default pattern to be used to add or delete index levels when generating the backup data set name from the name of the data set being reorganized. FDRREORG will use each index level specified in BACKUPINDEX in place of the original index level. If a period is specified without any characters following, the original index level will be copied to the backup data set name. If + is specified, the character following the + will be inserted into the backup data set name as a new index level. If ++ is specified, the characters following the ++ will be added to the end of the backup data set name as a new index level. If - is specified, the index level will be dropped from the backup data set name. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF userid or RACF groupid be used. See the description of the SELECT command in Section 30 for examples. The default is ++BACKUP.
LASTAPE- PREFIX	Specifies the high level qualifier to be used when cataloging or locating the special LASTAPE catalog entries. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF userid or RACF groupid be used. The default is FDRREORG.
BACKUP- FAILURE	Defines how FDRREORG should treat a backup failure that occurs as part of a reorganization. Backup failures that occur as a result of datasets selected via the NOREORG or ALWAYSBACKUP keywords will always be treated as an error. ERROR - The failure is treated as an error and will result in a return code of 8. WARNING - The failure is treated as a warning and will result in a return code of 4. The default is ERROR.
IAM- COMPPERM	If set to YES, FDRREORG will backup compressed IAM data sets in compressed format when BACKUP=GDG or BACKUP=PERM has been specified; this saves considerable time and space compared to decompressing the IAM data during the backup. If set to NO, compressed IAM data sets will be decompressed during output to permanent backups. The default is NO.

90.23 CONTINUED

**IAM-
COMPTMP** If set to YES, FDRREORG will backup compressed IAM data sets in compressed format when BACKUP=TEMP has been specified; this saves considerable time and space compared to decompressing the IAM data during the backup. If set to NO, compressed IAM data sets will be decompressed during output to temporary backups. The default is NO.

**PANEL
A.I.4.12.2
CHECKPOINT
DEFAULTS**

```
----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----
COMMAND ==>

                CHECKPOINT RELATED DEFAULTS FOR REORG AND RECOVER COMMANDS

CKPTALLOC      CHECKPOINT ALLOCATION METHOD  (SMS, UNIT)..... UNIT

CKPTDATACLASS  CHECKPOINT SMS DATA CLASS NAME.....
CKPTMGMTCLASS  CHECKPOINT SMS MANAGEMENT CLASS NAME.....
CKPTSTORCLASS  CHECKPOINT SMS STORAGE CLASS NAME.....

CKPTUNIT       CHECKPOINT UNIT NAME..... SYSDA

CKPTPREFIX     CHECKPOINT DSN PREFIX..... FDRREORG
```

CKPTALLOC Defines the allocation method to use when allocating checkpoint files.

SMS -- Allocate by SMS storage class.

UNIT -- Allocate by unit name.

The default is UNIT.

Note: If set to SMS, a valid SMS storage class name must also be specified for CKPTSTORCLASS. If set to UNIT, a valid unit name must also be specified for CKPTUNIT.

**CKPTDATA-
CLASS** Specifies the SMS data class to use for the checkpoint file when CKPTALLOC is SMS.

**CKPTMGMT-
CLASS** Specifies the SMS management class to use for the checkpoint file when CKPTALLOC is SMS.

**CKPTSTOR-
CLASS** Defines a valid SMS storage class which will be used to allocate SMS managed checkpoint files when CKPTALLOC is SMS.

CKPTUNIT Defines a valid unit name containing disk devices to be used for allocating the checkpoint file when CKPTALLOC is UNIT. It may be any value valid for UNIT= in JCL as long as it contains only disks. The default is SYSDA.

CKPTPREFIX Specifies the high level qualifier to use when constructing the name of the checkpoint file. To make it possible for the RECOVER command to find the checkpoint file for a failed REORG, the checkpoint data set is cataloged with a name of &CKPTPREFIX.REORGCKP.jobname.Dccyyddd.Thhmmss. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF userid or RACF groupid be used. The default is FDRREORG.

90.23 CONTINUED

PANEL
A.I.4.12.3
LOG DEFAULTS

```

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----
COMMAND ==>

          LOG FILE RELATED DEFAULTS FOR REORG AND RECOVER COMMANDS
LOGALLOC   LOG FILE ALLOCATION METHOD  (SMS, UNIT)..... UNIT

LOGDATACLASS LOG FILE SMS DATA CLASS NAME.....
LOGMGMTCLASS LOG FILE SMS MANAGEMENT CLASS NAME.....
LOGSTORCLASS LOG FILE SMS STORAGE CLASS NAME.....

LOGUNIT     LOG FILE UNIT NAME..... SYSDA

LOGPREFIX   LOG FILE DSN PREFIX..... FDRREORG

```

LOGALLOC Defines the allocation method to use when allocating the log file.

SMS -- Allocate by SMS storage class.

UNIT -- Allocate by unit name.

The default is UNIT.

Note: If set to SMS, a valid SMS storage class name must also be specified for LOGSTORCLASS. If set to UNIT, a valid unit name must also be specified for LOGUNIT.

LOGDATA-CLASS Specifies the SMS data class to use for the log file when LOGALLOC is SMS.

LOGMGMT-CLASS Specifies the SMS management class to use for the log file when LOGALLOC is SMS.

LOGSTOR-CLASS Defines a valid SMS storage class which will be used to allocate SMS managed log files when LOGALLOC is SMS.

LOGUNIT Defines a valid unit name containing disk devices to be used for allocating the log file when LOGALLOC is UNIT. It may be any value valid for UNIT= in JCL as long as it contains only disks. The default is SYSDA.

LOGPREFIX Specifies the high level qualifier to use when constructing the name of the log file. To make it possible for the RECOVER command to find the log file, the log data set is cataloged with a name of &LOGPREFIX.REORGLOG.jobname.Dccyyddd.Thhmmss. You can specify a special identifier of &RACFUID or &RACFGID to indicate that the RACF userid or RACF groupid be used. The default is FDRREORG.

90.23 CONTINUED

PANEL A.I4.12.4
MISCELLANEOUS
DEFAULTS

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----	
COMMAND ===>	
MISC. DEFAULTS FOR REORG, SIMULATE, AND RECOVER COMMANDS	
MAXENVERR	MAXIMUM NUMBER OF ENVIRONMENTAL ERRORS ALLOWED99
MAXSYSERR	MAXIMUM NUMBER OF SYSTEM ABENDS ALLOWED99
MAXENQ	MAXIMUM NUMBER OF OUTSTANDING ENQ'S ALLOWED FOR DSNRETRY999
MAXTASKS	MAXIMUM NUMBER OF CONCURRENT VOLUMES TO BE PROCESSED4
MSGLEVEL	LOWEST LEVEL MESSAGE TYPE TO PRINT (I, W, E)I
OWNERSTRING	OWNERID OF THE DATA COMPONENT OF A VSAM KSDSFDR\$
SUBTASKABEND	ACTION TO TAKE IF A SUBTASK ABENDS (CONT OR TERM)CONT
SELTERROR	SET NON-ZERO RETURN CODE IF NO DATA SETS ARE SELECTEDYES
MSGTIMESTAMP	TIMESTAMP MESSAGES WITH DATE, TIME AND SUBTASK IDYES
SORTRPT	SORT THE REORG REPORT IN DATA SET NAME AND VOLSER ORDERNO
LISTNOREORG	PRINT THE FDRREORG NOREORG LIST ON REORG AND SIMULATEYES

MAXENVERR Specifies the maximum number of environmental errors allowed (1 to 99999). When this limit is reached, all subtasks will terminate after processing the currently active data set; no additional REORG or SIMULATE commands will be processed. Environmental errors are any backup or re-load failures not caused by a system ABEND. Insufficient space, or target data set not available, are examples of environmental errors. The default is 99.

MAXSYSERR Specifies the maximum number of system ABENDs allowed (1 to 99999). When this limit is reached, all subtasks will terminate after processing the currently active data set. No additional REORG or SIMULATE commands will be processed. The default is 99.

MAXENQ Specifies the maximum number of outstanding ENQs allowed (1 to 999) for the DSNRETRY ENQ and WAIT options for each volume processor task. When this limit is reached, data sets will be added to task retry queue as if the RETRY option was specified. The ENQ will be issued when the outstanding ENQ count falls below this maximum. Once the volume processor task has completed processing all volumes and is waiting for data sets queued with the WAIT option, no additional ENQs will be issued. The default is 999.

MAXTASKS Specifies the maximum number of concurrent volumes (1 to 15) to be processed. Note that the actual number of subtasks possible is limited by the amount of available virtual storage. For SIMULATE commands with MAXTASKS=15, a region size of 2.5 megabytes should be sufficient. For REORG commands with MAXTASKS=15, a region size of 8.5 megabytes should be sufficient. The default is 4.

MSGLEVEL Specifies the lowest level message type to be displayed on the subtask print file which has the ddname REORGPRT.

I -- Informational, warning, and error messages will be displayed

W -- Warning and error messages will be displayed

E -- Only error messages will be displayed

The default is I.

OWNER-STRING Specifies a string of up to four characters which will be used as an eye catcher in the ownerid of the data component of a VSAM KSDS. FDRREORG will update the ownerid field of the data component's catalog entry to record information used to insure data set integrity and to identify data sets that were being processed by FDRREORG when either a system failure occurred, or FDRREORG was canceled or otherwise terminated. See the RECOVER command for additional information. The default is FDR\$.

90.23 CONTINUED

SUBTASK-ABEND Specifies the action to take if a volume processor subtask ABENDs.

CONT -- Continue processing without the subtask.

TERM -- Quiesce all active work and terminate.

The default is CONT.

SELTERROR If set to YES, FDRREORG will terminate with a return code of 8 if no data sets are selected by REORG or SIMULATE; this may help draw your attention to an error in your SELECT statements. If set to NO, the FDRREORG return code will be 0 if no data sets are selected (unless other errors occur); this can be used when the selection criteria you use may naturally result in no selections. The default is YES.

MSGTIME-STAMP If set to YES, messages written to REORGPRT will be suffixed with the current date, time, and the internal subtask id that issued the message; the LRECL of the REORGPRT file is increased to 151. The default is YES.

Note: Specify NO if you do not have a printer capable of printing 150 characters per line.

SORTRPT If set to YES, the report written to REORGRPT at the end of a REORG, SIMULATE or RECOVER function will be sorted in data set name and volume serial order; this requires that information related to each data set processed or selected be kept in storage until the function is completed. If set to NO, the report is produced as each data set is processed or selected. The default is NO.

LISTNOREORG If set to YES, the NOREORG list is printed for REORG and SIMULATE functions. If set to NO, the list is not printed. The default is YES.

90.23 CONTINUED

PANEL A.I.4.12.5
MORE
MISCELLANEOUS
DEFAULTS

```

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----
COMMAND ==>

                MORE MISC. DEFAULTS FOR REORG, SIMULATE, AND RECOVER COMMANDS

UPDATEDPDS      COMPRESS PDS'S WHICH DO NOT HAVE A CURRENT BACKUP ..... YES
NOUPDATES       REORGANIZE VSAM AND IAM FILES THAT HAVE NO UPDATES ..... NO
ALIASCHECK      CHECK FOR ALIAS PRIOR TO ALLOCATING A BACKUP DATA SET ..... YES

VSAMDEFINE      REDEFINE NON-REUSABLE VSAM KSDS   (NO, IFREQ, ALWAYS) ..... IFREQ

POOLDASDUNIT    POOLDASD IS INSTALLED IN THE SYSTEM ..... NO
KEEPCRDATE      RETAIN THE ORIGINAL CREATION DATE WHEN REDEFINING VSAM ..... YES
VSAMAUTOSPACE   INCREASE SPACE IF RELOAD FAILS DUE TO OUT OF SPACE ..... YES

REDUCEPRIMARY    REDUCE PRIMARY SPACE ALLOCATION TO ACTUAL SPACE USED ..... YES

ENQERR          TREAT DATA SET ENQ FAILURE AS AN ERROR (YES,NO,WARN) ..... YES

```

- UPDATEDPDS** If set to YES, partitioned data sets (PDSs) which do not have a current backup will be compressed (data sets are considered not to have a current backup if the UPDATE indicator is on in their Format 1 DSCB, indicating it has been updated since the last backup was taken). If set to NO, PDSs with no current backup are bypassed. The default is YES.
- NOUPDATES** If set to YES, IAM and VSAM data sets selected for reorganization that have had no adds, no deletes, and no updates (no update activity) will still be reorganized. If set to NO, such data sets are bypassed. The default is NO.
- ALIASCHECK** If set to YES, a check will be done to see if a catalog alias exists before allocating a backup data set, i.e., will the backup data set be cataloged in a user catalog or in the master catalog; if no alias exists, the backup or reorganization will not be performed. If set to NO, no alias check is done and the backup data set may be cataloged in the master catalog (if authorized to do so) if no alias exists. The default is YES.
- VSAMDEFINE** Specifies when FDRREORG should delete and define VSAM KSDSs that will be reorganized.
ALWAYS - VSAM KSDSs are always deleted and defined before reloading.
IFREQ - Only VSAM KSDSs that cannot be reused will be deleted and defined before reloading.
NO - Disables delete and define of all VSAM KSDSs. KSDSs that cannot be reused will not be reorganized.
The default is IFREQ.
- POOL DASDUNIT** If set to YES, unit names which are not defined to the system (not in the system EDT - Eligible Device Table) will be accepted for allocation of backup, recovery and log data sets. If set to NO, unit names must be in the EDT or the allocation will not be attempted. This should be set to YES if you have the POOL-DASD product from Boole and Babbage or similar DASD allocation software which accepts unit names not in the EDT. The default is NO.
- KEEPCRDATE** If set to YES, FDRREORG will retain the original creation date when redefining VSAM files. If set to NO, the redefined VSAM files will have the current date as the creation date. The default is YES.
- VSAM AUTOSPACE** If set to YES, FDRREORG will attempt to increase the space allocation of a single volume VSAM KSDS if a reload fails due to an out of space condition; it will calculate how much additional space is required to reload the data set, redefine it and retry the reload. If set to NO, out-of-space failures will not be retried. The default is YES.
- REDUCE-PRIMARY** If set to YES, FDRREORG will reduce the space allocation of VSAM KSDSs and Enhanced Format IAM files to the space actually used by the data. If set to NO, the original allocation is retained. The

90.23 CONTINUED

default is YES.

ENQERR Controls processing of ENQ failures on data sets:
YES – return code 8 is set if a required data set is in use
NO – return code 0 is set even if a required data set is in use.
WARN – return code 4 is set if a required data set is in use.

PANEL
A.I.4.12.6
SELECTION
DEFAULTS

----- FDR INSTALLATION -- SET FDRREORG OPTIONS -----	
COMMAND ==>	
DEFAULT SELECTION PARAMETERS FOR SELECT COMMAND	
VSAM FILES	
CASPLITR	VSAM KSDS MINIMUM RATIO OF CA SPLITS FOR EVERY 100 CAS 10
CISPLITR	VSAM KSDS MINIMUM NUMBER OF CI SPLITS FOR EVERY 100 CIS 10
FREESPACE	VSAM CA OR CI FREESPACE PERCENTAGE SELECTION CUT OFF... 50
IAM FILES	
OFULL	MINIMUM PERCENTAGE OF INDEPENDENT OVERFLOW USED..... 90
PCTTRECO	MINIMUM PERCENTAGE OF RECORDS IN INDEPENDENT OVERFLOW.. 10
PEFULL	MINIMUM PERCENTAGE OF PRIME EXTENTION USED..... 100
PEUDATAR	PCT OF USED PRIME EXTENTION BLOCKS TO PRIME DATA BLOCKS 100
OVERFLOWINDEX	AMOUNT OF MEMORY FOR THE IN-STORAGE IAM OVERFLOW INDEX. 1048576
PDS FILES	
PDSFULL	MINIMUM PERCENTAGE OF ALLOCATED SPACE USED BY A PDS.... 90
PDSEXTENTS	MINIMUM NUMBER OF EXTENTS FOR A PDS..... 17

CASPLITR Defines the default minimum ratio of Control Area splits to every 100 Control Areas in a VSAM KSDS. The default is 10.

CISPLITR Defines the default minimum ratio of Control Interval splits to every 100 Control Intervals in a VSAM KSDS. The default is 10.

FREESPACE Defines the default selection cut off for VSAM CA or CI free space percent. This value is used to prevent reorganizations of VSAM KSDSs or AIXs that have very high free space percentages. To reduce splits for some files in online systems, some VSAM files are loaded with little or no free space and then altered to have a very high CA and/or CI freespace percent. If these files are processed by FDRREORG, the current freespace percentage is used resulting in a dramatic increase in space. The default is 50.

OFULL Defines the default minimum percentage of independent overflow used in an IAM file. The default is 80.

PCTTRECO Defines the default minimum percentage of total records in an IAM file that are in the IAM independent overflow area. The default is 10.

PEFULL Defines the default minimum percentage of prime extension used in a IAM file. The default is 100.

PEUDATAR Defines the default minimum ratio times 100 of used prime extension blocks to prime data blocks in an IAM file. The default is 100.

OVERFLOW-INDEX Specifies the amount of memory in bytes required for the in-storage IAM overflow index. The default is 1048576 (1 megabyte).

PDSFULL Defines the default minimum percentage of allocated space used by a PDS. The default is 90.

PDSEXTENTS Defines the default minimum number of extents for a PDS. The default is 17.

90.24 ENABLE FDRCLONE OPTIONS (PANEL A.I.4.13)

If you are licensed for FDRCLONE, described in detail in [Section 50.70](#), you must enable FDRCLONE operation. However, **this is done only on the FDRCLONE target system**, the LPAR or disaster/recovery system on which the data sets are to be cloned. Once you install FDR on your clone target system (or restore the FDR libraries from your production system), you will enable FDRCLONE data set restores by setting options on this panel.

**PANEL A.I.4.13
FDRCLONE
OPTIONS**

```
----- FDR INSTALLATION -- SET FDRCLONE OPTIONS -----
COMMAND ==>

FDRCLONE  FDRCLONE RESTORE IS ENABLED..... NO
MAXTAPES  MAXIMUM NUMBER OF TAPE UNITS FOR FDRCLONE RESTORES..... 10
LEAVETAPE NUMBER OF MINUTES TO LEAVE A TAPE MOUNTED..... 0
```

- FDRCLONE** This should be set to **NO** on your normal production and test systems. Change it to **YES** on the FDRCLONE target system to enable FDRCLONE to restore required data sets dynamically on the clone system. [See Section 50.70](#) for complete details.
- MAXTAPES** This is the maximum number of tapes volumes which FDRCLONE can mount simultaneously while doing restores, from 1 to 10. If you need to reserve some tape drives for batch jobs, set this to a value lower than the number of drives on the clone target system.
- LEAVETAPE** This is the number of minutes that FDRCLONE will keep a tape mounted after all restores are complete, to see if another request for a restore from the same tape will be issued. Setting this to a non-zero value during the initial period of a FDRCLONE system recovery will minimize the number of times that tape must be mounted and will reconstruct the system more rapidly. As the frequency of FDRCLONE restore requests diminishes you may wish to set it to zero to dismount tapes immediately when they are no longer required.

90.30 DEFINE THE COMPAKTOR UNMOVABLE TABLE (PANEL A.I.5)

NEW INSTALLATIONS

If you have COMPAKTOR, you will want to review the default COMPAKTOR Unmovable Table and probably will update it. [See Section 40.21](#) for guidance on the data sets you may need to add to the table.

EXISTING INSTALLATIONS

If you have COMPAKTOR, you can use the COMPAKTOR Unmovable Table from your previous release. It can be copied in this dialog (the COPY option), or you may have already copied it with the COPY option on Panel A.I.4. In either case, you may want to review and update it.

Select option 5 (SETCPK) on the FDR Installation Options Menu which will bring you to this panel. Like other FDR options, the COMPAKTOR Unmovable Table resides in a load module in the FDR program library.

**PANEL A.I.5
COMPAKTOR
UNMOVABLE
TABLE**

```
----- FDR INSTALLATION -- COMPAKTOR UNMOVABLE TABLE -----
COMMAND ==>

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME      ==> 'IDP.MODFDR53'
  VOLUME SERIAL       ==>
```

**FDR PROGRAM
LIBRARY
DATA SET**

If the name of the FDR program library is not correct, correct it (optionally you can specify the volume serial if it is not cataloged for some reason). This is the library where the COMPAKTOR unmovable table will be stored; it must already contain a copy of the module FDRCPKUM.

**REFRESHING
THE TABLE IN
LLA**

If the FDR program library is in the system link list, even though you save the updated table, the LLA facility of the operating system may still have a back-level version of the FDRCPKUM module. However, if you execute FDRSTART (the FDR dynamic exit installation program), it will also refresh the LLA copies of all of the FDR lists and tables. FDRSTART can be executed using the REFRESH command on the FDR Installation Options Menu (Panel A.I.4, [see Section 90.11](#)).

90.30 CONTINUED

**COMPAKTOR
UNMOVABLE
TABLE PANEL**

Once you press ENTER on the previous panel, this panel is displayed, showing the current contents of the COMPAKTOR Unmovable Table and allowing you to modify it. The example shown below is the default Unmovable Table distributed by Innovation.

----- FDR INSTALLATION -- COMPAKTOR UNMOVABLE TABLE -----			Row 1 to 3 of 3
COMMAND ==>			SCROLL ==> PAGE
SAVE	-	SAVE CHANGES	COPY - COPY THE MODULE
RESET	-	RESET MODULE	FIND - FIND A STRING
			CANCEL - EXIT IMMEDIATELY
			HELP - TUTORIAL PANELS
	DSN/		
CMD	DSG	DATA SET NAME OR GROUP	
---	---	-----	
	DSG	SYS1.VTOCIX.	
	DSG	SYS1.VVDS.	
	DSN	SYS1.LOGREC	
***** Bottom of data *****			

The COMPAKTOR Unmovable Table specifies data sets that COMPAKTOR will leave in the same location when it reorganizes disk volumes. FASTCPK automatically recognizes data sets in use and makes them unmovable, but certain data sets are not properly ENQed by the system so COMPAKTOR cannot recognize them as active. This table may need to include your LINKLIST data sets, tape management data sets, JES data sets and the other data sets listed in [Section 40.21](#) that are not ENQed by the system.

The DSN/DSG column indicates whether the associated entry is a fully-qualified data set name (DSN) or a data set prefix (DSG). In either case the name may also be a data set name mask, as described in [Section 80](#). For example DSN **.LIST will make all data sets ending in LIST unmovable. DSG entries have an implied "***" at the end.

Basic line commands (I -- insert; D -- delete; R -- repeat) are used to add and delete entries in the table.

You may exclude (protect) certain volumes from COMPAKTOR processing by making special entries in the COMPAKTOR Unmovable Table. If COMPAKTOR finds a name in the Unmovable Table in the format:

FDRCPK . EXCLUDE . COMPAKT . Vvvvvvvv

any attempt to do any COMPAKTion on volume "vvvvvv" will result in a warning. If the name is:

FDRCPK . EXCLUDE . RELEASE . Vvvvvvvv

then space release (TYPE=RLSE) will be inhibited. If the name in the table is a DSN entry, then only the named volume is excluded. If the entry is a DSG, then "vvvvvv" must be 5 or fewer characters and all volumes starting with that prefix will be excluded.

For example, entries of

DSN FDRCPK . EXCLUDE . COMPAKT . VTS0123
DSG FDRCPK . EXCLUDE . RELEASE . VTS0

will exclude volume TS0123 from COMPAKTion functions, and will exclude all TSO volumes from release processing.

The source code for module FDRCPKUM, which contains the COMPAKTOR Unmovable Table, is supplied in the Installation Control Library. If the number of entries in the table is inadequate you may change the source code to allow for more entries and then assemble and linkedit the module.

On the COMMAND line, enter SAVE to save the updated table, RESET to reset it to the Innovation default table, and COPY to copy the table from a previous release.

90.31 DEFINE THE FDRREORG NOREORG LIST (PANEL A.I.5A)

NEW INSTALLATIONS

If you have FDRREORG, you may want to review the default FDRREORG NOREORG List and may want to update it, to inhibit certain data sets from ever being reorganized.

EXISTING INSTALLATIONS

If you have FDRREORG, you can use the FDRREORG NOREORG List from your previous release. It can be copied in this dialog (the COPY option), or you may have already copied it with the COPY option on Panel A.I.4. In either case, you may want to review and update it.

Select option 5A (SETREORG) on the FDR Installation Options Menu which will bring you to this panel. Like other FDR options, the FDRREORG NOREORG List resides in a load module in the FDR program library.

**PANEL A.I.5A
FDRREORG
NOREORG LIST**

```
----- FDR INSTALLATION -- FDRREORG NOREORG LIST -----
COMMAND ==>
PLEASE PRESS THE "ENTER" KEY TO DISPLAY THE TABLE

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME      ==> 'IDP.MODFDR53'
  VOLUME SERIAL       ==>
```

**FDR PROGRAM
LIBRARY DATA
SET**

If the name of the FDR program library is not correct, correct it (optionally you can specify the volume serial if it is not cataloged for some reason). This is the library where the NOREORG List will be stored; it must already contain a copy of the module FDRNORG.

**REFRESHING
THE LIST IN
LLA**

If the FDR program library is in the system link list, even though you save the updated list, the LLA facility of the operating system may still have a back-level version of the FDRNORG module. However, if you execute FDRSTART (the FDR dynamic exit installation program), it will also refresh the LLA copies of all of the FDR lists and tables. FDRSTART can be executed using the REFRESH command on the FDR Installation Options Menu (Panel A.I.4, [see Section 90.11](#)).

90.31 CONTINUED

**NOREORG LIST
PANEL**

Once you press ENTER on the previous panel, this panel is displayed, showing the current contents of the FDRREORG NOREORG List and allowing you to modify it. The example shown below is the default NOREORG List distributed by Innovation (all SYS1 data sets are inhibited).

Note: If you execute FDRCOPY directly to compress a PDS, the NOREORG list will not be used.

```

----- FDR INSTALLATION -- FDRREORG NOREORG LIST ----- Row 1 to 1 of 1
COMMAND ===>                                     SCROLL ===> PAGE

          SAVE - SAVE CHANGES    COPY - COPY THE MODULE    CANCEL - EXIT IMMEDIATELY
          RESET - RESET MODULE    FIND - FIND A STRING      . HELP  - TUTORIAL PANELS

CMD      NOREORG STATEMENT      (NOTE: VALIDATION IS DELAYED UNTIL MODULE IS SAVED)
---      ----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----6-----+-----7-
          NOREORG DSN=SYS1.**
          ***** Bottom of data *****
    
```

Basic line commands (**I** -- insert; **D** -- delete; **R** -- repeat) are used to add and delete entries in the list.

Each entry must begin with a NOREORG command as shown above. NOREORG supports 2 operands:

DSN= which defines one or more data set names or name masks (as described in [Section 80](#)). If multiple names are specified, enclose them in parenthesis, e.g., DSN=(SYS1.**;SYS2.**)

VOL= which defines one or more DASD volume serials or volume serial prefixes. A prefix is specified by following it with an asterisk, e.g., VOL=ABC*. If multiple serials or prefixes are specified, enclose them in parenthesis, e.g., VOL=(TSO123,ABC*)

If DSN= appears on a NOREORG statement without VOL=, the named data sets will be protected from reorganization when they are found on any DASD volume.

If VOL= appears on a NOREORG statement without DSN=, all data sets on the named volumes will be protected from reorganization.

If DSN= and VOL= both appear, the named data sets are protected only when they are found on one of the named volumes.

You may include multiple NOREORG statements in the list. If the DSN= and VOL= parameters are too long to fit on one line, you can continue the statement by ending the line after any command and inserting the continuation on the next line.

On the COMMAND line, enter SAVE to save the updated list, RESET to reset it to the default, and COPY to copy the list from a previous release.

90.32 DEFINE THE ABR PROTECT LISTS AND RESTORE ALLOCATION LIST (PANEL A.I.6)

NEW INSTALLATIONS

If you have ABR, review the functions of the ABR Protect Lists and Restore Allocation List described below. If necessary, define the lists as shown.

EXISTING ABR INSTALLATIONS

If you have ABR, the ABR lists can be copied in this dialog (the COPY option), or you may have already copied then with the COPY option on Panel A.I.4. In either case, you may want to review and update them.

Like other FDR options, the ABR lists are stored as load modules in the FDR program library; as distributed they can hold a limited number of entries (100 entries or statement lines). The source code for the Protect Lists and Restore Allocation Control List is supplied in the Installation Control Library. You may change the source code to allow more entries and then assemble and linkedit the module.

PANEL A.I.6
ABR PROTECT
LISTS AND
RESTORE
ALLOCATION
LIST

Select option 6 on the FDR Installation Options Menu which will bring you to this panel.

```

----- FDR INSTALLATION -- ABR PROTECT LISTS AND RESTORE ALLOCATION LIST --
OPTION  ===>                                SCROLL ===> PAGE

      1  ALLOCATE - ABR RESTORE ALLOCATION LIST... ACTIVE ===> NO
      2  ARCPROT  - ABR ARCHIVE PROTECT LIST..... ACTIVE ===> NO
      3  ABRPROT  - ABR BACKUP  PROTECT LIST..... ACTIVE ===> NO
      4  RESTPROT - ABR RESTORE PROTECT LIST..... ACTIVE ===> NO
      5  SCRPROT  - ABR SCRATCH PROTECT LIST..... ACTIVE ===> NO

FDR PROGRAM LIBRARY DATA SET:
DATA SET NAME      ===> 'IDP.MODFDR53'
VOLUME SERIAL      ===>
    
```

ABR LISTS

The 5 ABR lists (the Restore Allocation List and the 4 Protect Lists) are listed here. To use a given list, you must both activate it and define it. The lists are activated by changing the NO to the right of each list to a YES (these activation options are saved in FDROPT, the main FDR options module).

You must now define the lists that you have activated (or intend to activate later). The list numbers (1 through 5) must be entered one at a time next to OPTION on the top of the screen. For example, enter "1" to define the Restore Allocation List. Any list which is activated but not defined will be empty.

FDR PROGRAM
LIBRARY DATA
SET

If the name of the FDR program library is not correct, correct it (optionally you can specify the volume serial if it is not cataloged for some reason). This is the library where the ABR lists will be stored; it must already contain a copy of the module FDROPT (to store the activation options), and modules FDRALLOC, FDRPROTA, FDRPROTD, FDRPROTR, and FDRPROTS (depending on which lists you are updating).

REFRESHING
THE LISTS IN
LLA

If the FDR program library is in the system link list, even though you save the updated lists, the LLA facility of the operating system may still have a back-level version of the list modules. However, if you execute FDRSTART (the FDR dynamic exit installation program), it will also refresh the LLA copies of all of the FDR lists and tables. FDRSTART can be executed using the REFRESH command on the FDR Installation Options Menu (Panel A.I.4, [see Section 90.11](#)).

90.32 CONTINUED

PANEL A.I.6.1
RESTORE
ALLOCATION
LIST

The ABR Restore Allocation List gives you the capability to specify to ABR where specific data sets will be restored if they cannot be restored to their target volume (usually the original volume). This is used primarily with archived data sets, where the original volume may no longer exist or may have no room for the data set. The list distributed by Innovation is empty, but the sample list below has some entries filled in as examples. It is stored in module FDRALLOC in the FDR program library.

If the LXCONUSE option is set to YES and the LXCONVOL option is set to a non-existent volume serial on panel A.I.4.11.3 (section 90.22), then automatic recalls will not default to the original volume and will use the RESTORE Allocation List instead.

RESTORE
ALLOCATION
LIST EXAMPLE

```
----- FDR INSTALLATION -- RESTORE ALLOCATION LIST ----- Row 1 to 4 of 4
COMMAND ===>                                           SCROLL ==> PAGE

      SAVE  - SAVE CHANGES      COPY - COPY THE MODULE CANCEL - EXIT IMMEDIATELY
      RESET - RESET MODULE      FIND - FIND A STRING  HELP  - TUTORIAL PANELS
CMD  ALLOCATE STATEMENT (NOTE: VALIDATION IS DELAYED UNTIL MODULE IS SAVED)
---  -----1-----2-----3-----4-----5-----6-----7-----
      ALLOCATE ALLDSN,VOL=TEST23,NVOL=TEST25
      ALLOCATE ALLDSN,VOLG=PROD,NVOLG=PROD
      ALLOCATE DSN=SYS2.** ,VOL=( SYS123 ,SYS124 ) ,NVOL=SYS456
      ALLOCATE DSN=SYS3.** ,NVOLG=SYS*
      ALLOCATE ALLDSN,NVOL=(POOL1*,POOL2*)
***** Bottom of data *****
```

Basic line commands (**I** -- insert; **D** -- delete; **R** -- repeat) are used to add and delete entries in the list. Each entry must begin with an **ALLOCATE** command as shown above. Each **ALLOCATE** statement specifies the data sets to which it applies (by original name, original volume serial, or both) and then gives alternate volumes to which the data sets may be restored. **ALLOCATE** supports these operands:

DSN= specifies a fully-qualified data set name (up to 44 characters) or a filter to be used for generic data set selection, as described in Section 80. Note that this is matched against the **original** data set name; if you are restoring to a new name, it is not compared to the new name.

ALLDSN specifies that all data sets will be managed by this **ALLOCATE** statement (usually qualified by **VOL=** or **VOLG=**).

Either **DSN=** or **ALLDSN** must be present.

VOL= specifies one or more DASD volume serials (multiple serials are enclosed in parenthesis, separated by commas, e.g., "VOL=(ABC123,XYZ321)"). Only those data sets which were backed up from the volume(s) indicated will be managed by this **ALLOCATE** statement.

VOLG= specifies a DASD volume serial prefix, e.g., **VOLG=ABC** indicates all volumes beginning with ABC. Only those data sets which were backed up from a volume starting with that prefix will be managed by this **ALLOCATE** statement.

You can specify either **VOL=** or **VOLG=** or neither, but not both. Note that they specify the **original** volume serial, not the current target restore volume serial. If both are omitted, the **ALLOCATE** statement manages all data sets that match **DSN=** (or **ALLDSN**).

NVOL= specifies one or more DASD volume serials (multiple serials are enclosed in parenthesis, separated by commas), giving alternate output volumes as described below. **NVOL=** can also specify a volume prefix (e.g., **NVOL=ABC*** which is equivalent to **NVOLG=ABC**), multiple prefixes (e.g., **NVOL=(ABC*,XYZ*)**) or a combination (e.g., **NVOL=(ABC*,XYZ123)**)

NVOLG= specifies a DASD volume serial prefix, e.g., **VOLG=XYZ** indicates all volumes beginning with XYZ, giving alternate output volumes as described below. **NVOLG=ABC** is equivalent to **NVOL=ABC***.

You must specify either **NVOL=** or **NVOLG=**, not both.

On the **COMMAND** line, enter **SAVE** to save the updated list, **RESET** to reset it to an empty list, and **COPY** to copy the list from a previous release.

...

90.32 CONTINUED

**RESTORE
ALLOCATE
LIST
OPERATION**

Every restore selects a target volume for each data set being restored. The rules for target volume selection are detailed in [Sections 50](#) (for restore from backup) and 51 (for restore from archive); additional rules for auto-recalls are shown in 51.

But if the allocation of the output data set on the target volume fails for any reason, including volume not online, no space, security, etc., ABR uses the Restore Allocation List (if active and defined) to choose one or more alternate volumes. It scans the list looking for the first statement where the DSN= (or ALLDSN) matches the original data set name, and the VOL= or VOLG= (if present) matches the volume serial of the volume from which the data set was originally backed up. Once such a ALLOCATE statement is found, the NVOL= or NVOLG= operands are used to select an alternate output volume or volumes.

Note that the volumes specified by NVOL/NVOLG will be used by ABR in the order that they are found in system UCB chain (Unit Control Blocks), not in the order you specify them.

For a manual restore, if an NVOL= operand is specified on the SELECT statement, and the allocation fails on the first volume specified, an ALLOCATE statement will be chosen based on the original volume serial, not the volume on which the failure occurred, and the remaining volumes specified on the SELECT NVOL= and the volumes given by NVOL/NVOLG on the matching ALLOCATE statement will be merged in UCB order to create the list of additional target volumes.

For the example Restore Allocate List shown in the panel above, the operation of that list would be:

ALLOCATE ALLDSN,VOL=TEST23,NVOL=TEST25 - all data sets originally backed up from volume TEST23 will be allocated on TEST25. This might be used when volume TEST23 no longer exists, so that data sets archived from that volume can be recalled.

ALLOCATE ALLDSN,VOLG=PROD,NVOLG=PROD - all data sets originally backed up from any DASD volume starting with PROD will be restored to any volume starting with PROD. This allows the PROD volumes to be used as a pool, so that any data set from that pool can be restored to any volume in the pool.

ALLOCATE DSN=SYS2. , VOL=(SYS123,SYS124),NVOL=SYS456** - all SYS2 data sets backed up from volume SYS123 or SYS124 will be allocated on SYS456. This might be used when the SYS2 data sets have been moved to a new volume, so that any recalled data sets will go to the correct volume. You would need to keep the ALLOCATE statement in the Restore Allocate List as long as any data sets that were archived from the original volumes remain in the ARCHIVE Control File.

ALLOCATE DSN=SYS3. , NVOL=SYS*** - all SYS3 data sets will be restored to any volume beginning with SYS.

ALLOCATE ALLDSN,NVOL=(POOL1*,POOL2*)- this is a "catchall" which says that any data set which was not directed to a new volume by the preceding statements, be allocated to one of a set of volumes starting with POOL1 or POOL2. A similar statement should usually be last in the Restore Allocate List to avoid restore failures for data sets not covered by preceding ALLOCATE statements.

Note: because the ALLOCATE statements in the Restore Allocate List are scanned in order, in this example a SYS2 or SYS3 data set backed up from a PROD volume will be restored to a PROD volume, since the second ALLOCATE statement will control them.

90.32 CONTINUED

**RESTORE
ALLOCATE
LIST ERRORS**

If your Restore Allocate List does not seem to be working correctly, check these things:

- is the syntax and spelling correct on all statements?
- is the order of the ALLOCATE statements appropriate (as described above)?
- have you specified the original data set names and volume serials (not the current target names and serials)?
- is the Restore Allocate List active? (if so, a batch ABR restore from ARCHIVE will list the ALLOCATE statements right after your RESTORE statement).

**ABR PROTECT
LISTS**

ABR has 4 Protect Lists, stored as modules in the FDR program library:

- the ARCHIVE Protect List protects data sets from being archived (DUMP TYPE=ARC) stored as module FDRPROTA.
- the BACKUP Protect List protects data sets against incremental backup (DUMP TYPE=ABR, AUTO, or DSF) stored as module FDRPROTD. Note that it never protects against full-volume ABR backups, which always backup every data set on every volume processed.
- the RESTORE Protect List protects against all types of ABR restores, including restore from backup and archive, stored as module FDRPROTR. Such data sets can only be restored by a full-volume restore.
- the SCRATCH Protect List protects against ABR SUPERSCRATCH (DUMP TYPE=SCR), stored as module FDRPROTS.

Notes:

- 1) Entries in the Protect Lists **cannot be overridden**. They provide almost absolute protection against the indicated operations for the data sets in the list. The only way that you can override them is to temporarily change or deactivate them, run the ABR job and change them back, or point to a STEPLIB containing empty copies of the ABR modules containing the lists.
- 2) However, the Protect Lists (except the RESTORE Protect List) are ignored when processing SMS-managed volumes with the SMSMANAGE=YES operand specified. In that case the SMS management class definition is used exclusively to select data sets for processing, ignoring the Protect Lists.

90.32 CONTINUED

The format of all 4 lists is identical, as is the format of the ISPF panels used to maintain them, so only the ARCHIVE protect list panel is shown as an example. The lists distributed by Innovation are empty, but the sample list below has a few entries as examples.

PANEL
A.I.4.6.2
ARCHIVE
PROTECT LIST

----- FDR INSTALLATION -- ARCHIVE PROTECTION LIST -----				Row 1 to 2 of 2	
COMMAND ==>				SCROLL ==> PAGE	
SAVE		- SAVE CHANGES	COPY - COPY THE MODULE	CANCEL - EXIT IMMEDIATELY	
RESET		- RESET MODULE	FIND - FIND A STRING	HELP - TUTORIAL PANELS	
CMD	DSN/			VOL/	VOLUME
---	---	-----		---	---
	DSN	GLOBAL.PARMLIB			
---	---	-----		---	---
	DSG	TCPIP		VOLG	SYS
---	---	-----		---	---
***** Bottom of data *****					

Basic line commands (**I** -- insert; **D** -- delete; **R** -- repeat) are used to add and delete entries in the list.

On each line, indicate whether the data set name is a fully-qualified name (DSN) or a data set name prefix (DSG), then enter the data set name or prefix. In either case the name may also be a data set name mask, as described in [Section 80](#). For example "DSN **.LIST" will protect all data sets ending in LIST. DSG entries have an implied "***" at the end.

You can optionally specify a volume serial (VOL) or volume serial prefix (VOLG); only those data sets matching the data set name on the volumes named will be protected. If omitted, those data sets are protected on all volumes on which they are found.

You can protect all data sets on a given volume or group of volumes by specifying "ALLDSN" for the data set name (the DSN/DSG column must be blank). A volume serial or volume serial prefix should be used with ALLDSN (otherwise all data sets in your installation would be protected, which rarely makes sense).

On the COMMAND line, enter SAVE to save the updated list, RESET to reset it to an empty list, and COPY to copy the list from a previous release.

Note that for the ARCHIVE and SUPERSCRATCH Protect lists, the ARCHIVE and SUPERSCRATCH functions can be disabled for entire volumes by setting a simple flag in the ABR Model DSCB on each volume (panel A.I.8), which is a superior way of protecting volumes against those functions.

PROTECT LIST
ERRORS

If one or more of your Protect Lists does not seem to be working correctly, check these things:

- is the syntax and spelling correct on all statements?
- is the Protect List active? (if so, a batch ABR job of the appropriate type will list the Protect List contents right after your DUMP or RESTORE statement).

90.40 AUTHORIZE THE FDR PROGRAMS

**ADD THE FDR
LIBRARY TO
AUTHORIZED
LIST**

The FDR load library must be defined to MVS as an APF authorized library. For ABR users, the installation steps in the following sections will not work until the library is authorized. If it is not already defined as authorized, you (or your MVS system programmer) must do so by:

- For MVS releases prior to ESA V4, update the IEAAPFxx member in PARMLIB. An IPL will be required to activate the updated authorized library list.
- For ESA V4 and above (including OS/390), if your installation has converted to the use of the PROGxx member in PARMLIB to define authorized libraries, you can update the PROGxx member and activate the updated list with the MVS console command:
`SET PROG=xx`

If your installation has specified that the authorized program list is in dynamic format, you can also temporarily authorize the FDR load library with the MVS console command:

`SETPROG APF,ADD,DSNAME=fdr.load.library,VOL=volser`

- You may also have another means of temporarily authorizing the FDR load library for installation and testing.

Remember that temporary authorization will vanish after the next IPL unless the appropriate PARMLIB member is updated.

COMPAKTOR and ABR customers may need to provide for authorized program execution under TSO.

**TSO
FOREGROUND
AUTHORIZA-
TION**

Some of the foreground options available through the FDR ISPF dialogs require that some FDR programs be listed in the TSO Authorized Program Table in order to be executed with APF (Authorized Program Facility) authorization.

**FUNCTIONS
THAT NEED
TSO
AUTHORIZA-
TION**

The functions that require APF authorization when performed in the foreground are FDREPORT (panel A.1.6), RESTORE (panel A.2), ARCHIVE and SUPERSCRATCH (panel A.3), BACKUP (panel A.4), and COMPAKTOR MAP or SIMULATE (panel A.C). All of these functions can also be performed in the background, or (except COMPAKTOR) via the remote queue; if you specify the ABRGLOBL variable FOREGRND(NO), these functions will not be available in the foreground, and you will not have to provide APF authorization, for programs FDRABR and FDRCPK.

ARCHIVE or BACKUP via the Remote Queue (A.3, A.4, A.5) requires APF authorization for program FDRABRUT if the ABRGLOBL variable DISKUPDT is specified or defaulted as YES, but not if it is specified as NO.

Additionally, the Search, Report and Services SRS Dialog (A.S) requires APF authorization for program FDRSRSA.

**TSO
AUTHORIZED
PROGRAM
TABLE**

Starting with TSO/E version 1.4, the list of authorized programs may be specified in member IKJTSOxx in SYS1.PARMLIB. For prior versions of TSO/E, consult your current ISPF Customization documentation for instructions.

Add the FDR program names (FDRABR, FDRCPK, FDREPORT, FDRSRSA and FDRABRUT) to both AUTHPGM and AUTHTSF NAMES lists in a IKJTSOxx member. Consult the "Initialization and Tuning" reference for your level of MVS for the proper syntax.

90.40 CONTINUED

**TSO PARMLIB
COMMAND**

The TSO PARMLIB command, available in TSO/E version 2 and higher, can list the IKJTSOxx member current in use as well as its contents. It can also syntax check a new IKJTSOxx member and to active a new member.

PARMLIB

will display IKJTSOxx member currently in use and **all** of as its contents

PARMLIB LIST(AUTHPGM)

or

PARMLIB LIST(AUTHTSF)

will display IKJTSOxx member currently in use and the named authorization list.

PARMLIB CHECK(xx)

will syntax check a new or updated IKJTSOxx member.

PARMLIB UPDATE(xx)

will activate member IKJTSOxx.

90.41 CREATE THE ABR CATALOG (PANEL A.I.7)

NEW INSTALLATIONS: If you have ABR, you must create the ABR catalog.

EXISTING ABR INSTALLATIONS: you should use the existing ABR catalog.

THE ABR CATALOG

ABR needs to save pointers to many of the backups it creates. Rather than creating and managing an external data base for this information, ABR uses MVS catalog management to store this information in a standard ICF user catalog. Also, ABR needs to save backup information for data sets which have been scratched from their DASD volumes.

Each type of entry must be in a single ICF catalog with a one-level alias; ABR does not support multi-level aliases for the ABR catalog. The two types of ABR catalog entries can be in 2 separate catalogs, but most installations put them in a common ABR catalog. You can even put other types of catalog entries in the same catalog or assign the ABR aliases to an existing catalog, but **these are strongly discouraged**.

If the ABR catalog must ever be expanded or moved, member MOVECAT in the FDR ICL (Installation Control Library) contains instructions and guidance.

BACKUP SECTION

The backup section of the ABR catalog contains an entry for every full-volume, incremental, and selective backup file created by ABR (files created by running FDRABR with DUMP TYPE=FDR, AUTO, ABR or DSF - [Section 50](#)). These entries normally begin with "FDRABR.". ABR automatically maintains this section, deleting old generations as new generations are created, retaining the number of generations specified for each volume ([see Section 90.42](#)) so backup entries are usually retained only for a month or so.

SCRATCH SECTION

The scratch section of the ABR catalog contains an entry for data sets that were deleted from disk after having been backed up by ABR (full-volume or incremental backup). These entries are only created if the ABR DADSM Preprocessing Exit is installed ([Section 90.60](#)). The exit prefixes a high-level index of '#' to the name of the data set being scratched, and catalogs the resulting name in the ABR catalog. These entries identify the backup file if a user requests ABR to restore a scratched data set. The PURGE SCRATCH command of program FDRABRCM ([Section 50.50](#)) will delete these entries when the backups with which they are associated are no longer in the catalog; this must be run periodically to keep the catalog from filling up.

ARCHIVE SECTION

The archive section of the ABR catalog contains entries for certain ARCHIVE backups (TYPE=ARC - [Section 51](#)) and sometimes application backups (DUMP TYPE=APPL - [Section 52](#)). These entries normally begin with "FDRABR.". ARCHIVE normally catalogs only the first backup created in each step and any backup that starts a new reel of tape; application backups normally do not catalog any backups. But if requested (ARCCAT=ALL or EXPDT=99000 in the JCL) ABR will catalog **all** those backups. Backups on disk are always cataloged. These backups are also recorded in an ARCHIVE Control File but the catalog entries may also be used for restores. The REORG command of program FDRARCH ([Section 51.40](#)) removes these entries when the reorganization process finds that the archive tape has expired.

Note: because ARCHIVE backups may be retained for years, when ARCHIVE backups are cataloged the catalog entries may also be present for years. If you have requested cataloging of all ARCHIVE backups, you may need to adjust the size of the ABR catalog to accommodate them.

CATALOG ALIASES

The ABR catalog must be assigned catalog aliases (in the MVS master catalog) of:

FDRABR - for the backup and archive entries

- for the scratch entries

If you have changed the high-level indexes used for these entries in the FDR Global Option Table ([see Section 90.16](#) - Panel A.I.4.5), use the values you have assigned instead.

90.41 CONTINUED

**CREATING THE
ABR CATALOG**

Select option 7 (ABRCAT) on the FDR Installation Options Menu which will bring you to this panel, which will create IDCAMS batch JCL for the allocation of the catalog.

```

----- FDR INSTALLATION -- CREATE THE ABR CATALOG -----
COMMAND ==>

      EDIT - EDIT CREATE JCL      SUBMIT - SUBMIT JOBCANCEL - EXIT IMMEDIATELY

ABR CATALOG NAME ==> CATALOG.ABRBASE
MASTER CATALOG   ==>                                     (OPTIONAL)

CATALOG TYPE      ==> ICF      ALLOCATION UNIT ==> CYL
SHARED CATALOG    ==> NO      PRIMARY QUANTITY ==> 5
CATALOG VOLUME    ==> SYS003 SECONDARY QUAN  ==> 2
DEVICE TYPE       ==> 3390    FREESPACE CI%    ==> 10  CA% ==> 20

HIGH LEVEL INDEX FOR ABR BACKUPS      ==> FDRABR(BLANK IF NOT DESIRED)
HIGH LEVEL INDEX FOR SCRATCH ENTRIES  ==> #(BLANK IF NOT DESIRED)

SYSOUT CLASS      ==> *
JOB STATEMENT INFORMATION:
==> //useridA JOB (ACCOUNT),'NAME',CLASS=A,
==> //          NOTIFY=userid
==> //*
==> //*
```

**CATALOG
NAME**

Specify the name you want to assign to the ABR catalog. The Innovation recommendation "CATALOG.ABRBASE" is shown as a default but you may change it to anything that meets your installation's requirements.

The master catalog name can be omitted unless the ABR catalog must be defined in the master catalog of a system which is not the system the job will execute on. The master catalog named must be defined as a user catalog on the system where the job will run.

**CATALOG
SPECIFICA-
TIONS**

This section describes the characteristics of the catalog to be allocated.

If the ABR catalog will be shared by multiple MVS systems, it must be defined in the master catalog of each of those systems, so specify YES for SHARED CATALOG. This will take you to another panel where you provide the master catalog name for each of those systems (each of them must be defined as a user catalog on the system where the job will run).

The remainder of the parameters specify where to allocate the catalog and how big to make it. According to estimates in the IBM manual "Managing Catalogs", about 200 bytes are required for a non-VSAM entry, which translates to about 3000 entries in a 3380 cylinder and about 3500 in a 3390 cylinder. From the descriptions earlier in this section you can estimate how many entries will exist in the catalog, but Innovation suggests that you generously overallocate the catalog to allow for expansion and free space. An allocation of 25 to 50 cylinders will be adequate in most installations.

ABR INDEXES

These fields will be filled in with the values for the ABRINDEX and SCRINDEX currently in your FDR Global Option Table; there should be no reason to override them.

**SUBMIT THE
JOB**

Update the job information at the bottom of the screen and then enter, on the COMMAND line at the top, either SUBMIT to immediately submit the job or EDIT to review, edit and submit it manually (or save it).

90.42 SET ABR DISK VOLUME PROCESSING OPTIONS (PANEL A.I.8)

NEW INSTALLATIONS: If you are installing ABR, you will need to set the ABR disk volume processing options for any disk volume on which you want to use ABR for backups or ARCHIVING (this is also referred to as "initializing the ABR volumes"). This process will create a MODEL DSCB (a data set with no space assigned) on the specified volumes; this ABR MODEL DSCB contains the options you assign and will control the ABR processing for the volume. In addition, ABR will use two reserved bytes (at decimal displacements 103 and 104) in the DSCB for each data set on the volume.

EXISTING ABR INSTALLATIONS: your ABR volume options should already be set, but you may need to use this dialog to change options or initialize new volumes.

Note: Setting ABR disk volume processing options on a volume **does not affect** any other data sets currently on the volume. This is not the same as "initializing" a volume with ICKDSF; it does not disturb any existing data.

This panel may also be used to display information about ABR options and current ABR backups for previously initialized volumes, and to change the ABR volume processing options.

Select option 8 (ABRVOL) on the FDR Installation Options Menu which will bring you to this panel

**PANEL A.I.8
SET ABR DISK
VOLUME
PROCESSING
OPTIONS**

```

----- FDR INSTALLATION -- SET ABR DISK VOLUME PROCESSING OPTIONS -----
COMMAND ==>
PLEASE PRESS THE "ENTER" KEY TO DISPLAY THE TABLE                                SCROLL ==> PAGE

THE ABR DISK VOLUME PROCESSING OPTIONS MUST BE SET PRIOR TO ABR EXECUTION.

SETTING THE ABR PROCESSING OPTIONS DOES NOT AFFECT NORMAL USE OF THE VOLUME.

FDR PROGRAM LIBRARY DATA SET:
  DATA SET NAME      ==> 'IDP.MODFDR53'
  VOLUME SERIAL       ==>

SYSOUT CLASS          ==> *

JOB STATEMENT INFORMATION:
  ==> //useridA  JOB (ACCOUNT),'NAME',CLASS=A,
  ==> //          NOTIFY=userid
  ==> // *
                                ==> /*

```

**FDR PROGRAM
LIBRARY DATA
SET**

If the name of the FDR program library is not correct, correct it (optionally you can specify the volume serial if it is not cataloged for some reason).

**JCL
INFORMATION**

Since batch jobstreams will be generated if any volumes must be initialized or changed, specify SYSOUT CLASS and JOB STATEMENT information appropriate for running the job.

```

----- FDR INSTALLATION -- SET ABR VOLUME PROCESSING OPTIO Row 1 to 1 of 1
COMMAND ==>                                     SCROLL ==> PAGE
PLEASE ENTER THE VOLUME SERIAL NUMBERS TO DISPLAY/SET ABR PROCESSING OPTIONS
  SUBMIT - SUBMIT FDRABRM BATCH JOB      EDIT   - EDIT FDRABRM BATCH JOB
  REFRESH - REFRESH ABR VOLUME INFO      FIND    - FIND THE SPECIFIED STRING

```

-----CURRENT BACKUP-----											
CMD	VOLUME	GEN	CYC	EXPD	AC	MOD	MAX GEN	RETPD	ARCHIVE	STORCLAS	FORCE
	SERIAL	DATE		EXPD2			MAX AC	RETP2	SCRATCH	ARCHI	
	NUMBER								OLDBKUP	ARCLW	
----	-----	---	----	-----	----	----	-----	-----	-----	-----	-----
----	----					NO	4	60	YES		NO
							10	0	NO	80	
									NO	50	
----	-----	---	----	-----	----	----	-----	-----	-----	-----	-----

To display one or more volumes, enter a volume serial or mask in the VOLUME SERIAL NUMBER field. You can enter:

- a complete serial number (1-6 characters)
- a serial number prefix, followed by an asterisk, e.g., ABC*, which will display all volumes starting with that prefix.
- a serial number mask, with one or more asterisks in positions in the serial; the asterisk will match any character, e.g., A*B*01
- a combination, such as A*B* which will select any volume with A in position 1 and B in position 3.
- a single asterisk (*) will select all online disk volumes in your installation (use this option with caution if you have many DASD volumes online).

Basic line commands (**I** -- insert; **D** -- delete; **R** -- repeat) are used to add and delete entries in the list. Inserted lines will have a blank volume serial which you can overwrite to add one or more volumes to the display. You can also overwrite the volume serial of an existing entry; it will be deleted and replaced with the volume(s) selected.

90.42 CONTINUED

**INITIALIZING
THE VOLUMES
(continued)**

The following example display was created by entering a volume serial of TSO* in an installation where ABR has been in use for a while, but where a new TSO volume has been recently added.

----- FDR INSTALLATION -- SET ABR VOLUME PROCESSING OPTIO										
Row 1 to 4 of 4										
COMMAND ===>										
SCROLL ==> PAGE										
SUBMIT - SUBMIT FDRABRM BATCH JOB										
EDIT - EDIT FDRABRM BATCH JOB										
REFRESH - REFRESH ABR VOLUME INFO										
FIND - FIND THE SPECIFIED STRING										
-----CURRENT BACKUP-----										
CMD	VOLUME	GEN	CYC	EXPDT	AC	MOD	MAX GEN	RETPD	ARCHIVE	STORCLAS
	SERIAL	DATE	EXPDT	EXPDT2			MAX AC	RETP2	SCRATCH	ARCHI

---	TSO001	215	1	2000/030	1	NO	5	35	YES	
		1999/360					50	0	NO	NO

---	TSO002	9	1	2000/030	1	NO	5	35	NO	
		1999/360	1999/365				50	5	NO	NO

---	TSO003	****	NO	MODEL	**	NO	4	60	YES	
							10	0	NO	NO

---									YES	
									NO	80

---									NO	50
									NO	50

**CURRENT
BACKUP**

For a volume already initialized for ABR, these fields show the most recent backup of the volume.

The GEN/CYC/DATE column shows the ABR generation, cycle and creation date of the most recent backup.

The EXPDT/EXPDT2 column shows the expiration date of the most recent full-volume backup taken (if a COPY2 backup was created, its expiration is also shown).

The AC column shows the number of auto-cycles taken in the current generation (significant only if you are doing TYPE=AUTO backups).

The MOD field is normally NO, which prevents you from modifying other CURRENT BACKUP fields. In the rare case where you need to correct information recorded about the most recent backup, you change it to YES and modify the other fields.

If the display contains "NO MODEL", the volume is not currently initialized for ABR processing. The rest of the fields on the line will usually contain the defaults from the FDR Global Option Table. However, if you display an initialized volume, then overtype its volume serial with that of an uninitialized volume, the values from the first volume will be retained and can be used to initialize the new volume, effectively copying the values.

For a volume which has been initialized but never backed up by ABR, the display will say "NO BACKUP".

**MAX GEN
MAX AC**

MAX GEN is the maximum number of generations (full-volume backups) that ABR will retain for this volume. As new generations are created, older generations, based on MAXGEN, are uncataloged from the ABR catalog.

MAX AC is the maximum number of auto-cycles permitted before a full-volume backup is forced. This is significant only if you intend to do TYPE=AUTO backups as described in [Section 50](#).

**RETPD
RETP2**

RETPD is the default retention period which will be used for COPY1 full-volume backups, if no other retention period or expiration date is provided at execution time (in JCL or on the ABR DUMP statement). RETP2 is the default retention for COPY2 full-volume backups; a value of 0(zero) for RETP2 cause the value of RETPD to be used for COPY2.

90.42 CONTINUED

**ARCHIVE
SCRATCH
OLDBKUP**

These enable ABR processing options for each volume. ARCHIVE enables ARCHIVE backups (DUMP TYPE=ARC) and SCRATCH enables SUPERSCRATCH (DUMP TYPE=SCR); since both of these will delete data sets from the volume, you may choose to disable these functions on some or all volumes.

OLDBKUP enables recording of old backups in the Format 1 DSCB of each data set. If disabled, only the most recent (current) backup is recorded; older backups may still be available but they are not recorded by data set and cannot be easily displayed and selected. If enabled, up to 13 old backups (14 total) will be recorded and can be displayed and selected for restore. [See Section 90.18](#) for more information.

By default, ARCHIVE is enabled, SCRATCH and OLDBKUP are disabled. **Innovation strongly recommends enabling OLDBKUP on all volumes.**

Warning: SCRATCH should be enabled only for volumes where you intend to run SUPERSCRATCH to delete certain data sets with no backup (no chance of recovery). For example, you might run SUPERSCRATCH on work (storage) volumes to delete obsolete temporary data sets.

STORCLAS

If the volume is SMS-managed, you must specify a valid SMS storage class name in order to initialize the volume for ABR processing. Any valid storage class can be used, it doesn't have to result in an allocation on the volume. For further details, see "SMS Considerations" in [section 50.40](#) and "ABR Volume Selection" in [section 70.12](#).

**ARCHI
ARCLOW**

These specify high and low thresholds for ARCHIVE processing. They are used only if THRESHOLD= is specified on the DUMP TYPE=ARC statement, as described in [Section 51](#).

FORCE

Change this to YES only if an initialization of a volume for ABR processing fails because of non-zero values in the reserved DSCB fields used by ABR (relative decimal displacements 103 and 104). FORCE will force the initialization to continue, zeroing those reserved bytes in all Format 1 DSCBs on the volume. You should verify that the bytes are not being used by some other product. Existing ABR customers can use FORCE on new volumes without concern **but if you are a new ABR customers or are trialing (testing) ABR and get the message indicating that the reserved bytes contain data, contact Innovation for assistance before proceeding.**

**PROCESSING
THE REQUESTS**

Once you have made any required changes to the parameter displayed, enter SUBMIT or EDIT on the COMMAND line to submit or edit the batch job which will be generated.

The panel will generate a jobstream to execute the FDRABRM utility ([see Section 50.40](#)) for:

- any displayed volume which shows "NO MODEL", an ABRINIT statement will be generated to initialize it for ABR processing by creating an ABR MODEL DSCB, using the values shown on the line for that volume.
- any displayed initialized volume where some fields on the right of the display have been changed, a MAINT statement will be generated to change the ABR MODEL DSCB.
- any displayed initialized volume where some fields under CURRENT BACKUP have been changed (if MOD is set to YES), a REMODEL statement will be generated to rebuild the ABR MODEL DSCB with the values shown.

Any previously-initialized volume for which no fields were changed will be ignored. You can bypass initializing some volumes by deleting them from the display (the D line command).

Once the FDRABRM batch job completes, you can refresh the ABR volume information on the ISPF panel by using the REFRESH command.

90.42 CONTINUED

**QUICK
INITIALIZA-
TION OF
VOLUMES**

To quickly initialize for ABR processing **all** online volumes which are not already initialized, use this procedure:

- display the initial (empty) panel. Change any defaults on the screen that are required (such as YES for OLDBKUP).
- Under VOLUME SERIAL NUMBER put an asterisk (*) to select all online volumes.
- Enter EDIT or SUBMIT to generate ABRINIT statements, using the defaults you entered, for all uninitialized volumes.

90.43 CREATE THE ARCHIVE CONTROL FILE (PANEL A.I.9)

NEW INSTALLATIONS: If you are installing ABR and intend to do ARCHIVING, you must create at least one common ARCHIVE Control file.

EXISTING ABR INSTALLATIONS: Use your existing ARCHIVE Control File.

**THE ARCHIVE
CONTROL FILE**

When executing the ARCHIVE-option of ABR (DUMP TYPE=ARC), ABR can BACKUP, SCRATCH, and UNCATALOG data sets that match user selection criteria. When a data set is removed from the disk volumes, ABR must note the volume it came from, where the backup files were created, the DCB attributes, and the space the data set occupied. ABR uses the ARCHIVE Control File to contain pointers to the ARCHIVED data sets.

The ARCHIVE Control File must reside on a disk volume which has been initialized for ABR processing as described in the previous section; its data set name must include an index level starting with 'ARCHIVE' as any level other than the high-level index. The remaining characters are up to you.

There may be many Archive Control Files in an installation. Generally there is one common ARCHIVE Control File for normal ARCHIVING and AUTO-RECALL; if ARCHIVING or AUTO-RECALL are done from more than one CPU in an installation, this Control File should be shared across all the CPUs. However, there may be other ARCHIVE Control Files for special purposes.

**AUTO RECALL
AND
DYNAMIC
ALLOCATION**

The FDR Global Option table contains the name of the ARCHIVE Control File which will be dynamically allocated by the DYNARC option during ARCHIVE dumps and restores. This was set on Panel A.I.4.5 (see [Section 90.16](#)). This ARCHIVE Control File name is used for AUTO-RECALL by the Catalog Locate Exit and Data Set Not Found Exit, by the ABR ISPF panels, and by any ARCHIVE DUMP or RESTORE which specifies DYNARC. Since there can only be one such file named, one common ARCHIVE Control File should be used for all normal ARCHIVED data sets including all those set up for AUTO-RECALL (RECALL=YES specified or defaulted).

The default name is FDRABR.ARCHIVE. If you choose a different name, be sure to change it on Panel A.I.4.5; the name chosen must contain an index level of "ARCHIVE" anywhere but the first level.

**ARCHIVE
CONTROL FILE
ALLOCATION**

The ABR utility program, FDRARCH, will initialize the ARCHIVE Control File; it is described in [Section 51.40](#). The ISPF dialog in this section will generate an FDRARCH jobstream to initialize a Control File. Approximate space requirements for the ARCHIVE Control File at default blocking (2/track) are as follows:

3380 disk - 292 data sets/track

3390 disk - 348 data sets/track

The above applies to non-VSAM data sets; allow 4 data set entries for each VSAM cluster. Also multi-volume data sets will get one entry per volume.

Each entry in the ARCHIVE Control File will be retained for 1 year, by default. If the default retention period is used, there should be enough space allocated to the ARCHIVE Control File to accommodate 1 year of ARCHIVE activity. You must estimate the ARCHIVE activity within your own environment (the FDRQUERY program, described in [Section 10](#), may help).

Note: you will need to periodically reorganize all ARCHIVE Control Files, using the REORG function of FDRARCH ([Section 51.40](#)), in order to eliminate expired and obsolete entries.

90.43 CONTINUED

**PANEL A.I.9
ABR ARCHIVE
FILE
INITIALIZA-
TION**

Select option 9 (BLDARC) on the FDR Installation Options Menu which will bring you to this panel

```

----- FDR INSTALLATION -- ABR ARCHIVE FILE INITIALIZATION -----
COMMAND ==>

  EDIT - EDIT ARCINIT JCL      SUBMIT - SUBMIT JOB      CANCEL - EXIT IMMEDIATELY

  ARCHIVE DATA SET NAME      ==> FDRABR.ARCHIVE
  STANDARD DATA SET NAME     ==> YES
  MAX DATA SET ENTRIES       ==> 10000
  BLOCKS PER TRACK            ==> 2

  VOLUME SERIAL NUMBER        ==>

  FDR PROGRAM LIBRARY NAME     ==> 'IDP.MODFDR53'
  VOLUME SERIAL                ==>

  SYSPRINT SYSOUT CLASS       ==> *

  JOB STATEMENT INFORMATION:
  ==> //useridA JOB (ACCOUNT),'NAME'
  ==> /*
  ==> /*
  ==> /*

```

**ARCHIVE
CONTROL FILE
PARAMETERS**

ARCHIVE DATA SET NAME: Specify the fully-qualified data set name of the ARCHIVE Control File to be created.

STANDARD DATA SET NAME: If set to YES, the data set name must contain an index level of "ARCHIVE". Innovation recommends that you do not change this.

MAX DATA SET ENTRIES: Specify the maximum number of data set entries that you expect will need to be stored in the Control File, up to 3 million (3000000). This will be used to calculate the disk space required.

BLOCKS PER TRACK: Specify the blocking factor (from 1 to 5). 2 is recommended.

VOLUME SERIAL NUMBER: Specify the DASD volume where the Control File is to be allocated. Direct SMS allocation is not supported, but if SMS directs the allocation to a SMS-managed volume this will work.

**FDR PROGRAM
LIBRARY
DATA SET**

If the name of the FDR program library is not correct, correct it (optionally you can specify the volume serial if it is not cataloged for some reason).

**JCL
INFORMATION**

Since a batch jobstream will be generated to execute FDRARCH, specify SYSOUT CLASS and JOB STATEMENT information appropriate for running the job.

**SUBMITTING
THE JOB**

Once you have made any required changes to the parameters displayed, enter SUBMIT or EDIT on the COMMAND line to submit or edit the batch job which will be generated.

90.44 SET FDR DIALOG GLOBAL OPTIONS (PANEL A.I.10)

NEW INSTALLATIONS: You must set the options which will be used by the rest of the FDR ISPF dialogs.

EXISTING INSTALLATIONS: You should use the COPY function to copy dialog options from the CLIST library of your previous FDR release.

**PANEL A.I.10
SET FDR
DIALOG GLOBAL
OPTIONS**

Select option 10 (DIALOG) on the FDR Installation Options Menu which will bring you to this panel

```
----- FDR INSTALLATION -- SET FDR DIALOG GLOBAL OPTIONS -----
COMMAND ==>

THE FDR DIALOG CLIST LIBRARY CONTAINS A MEMBER ABRGLOBL WHICH MUST BE TAILORED
TO REFLECT THE USER'S ENVIRONMENT.

THE ABRGLOBL CLIST CONSISTS OF GLOBAL VARIABLES THAT ARE REQUIRED FOR EXECUTION
OF THE ABR ISPF/PDF COMMANDS.

THE DESCRIPTION OF ALL ABRGLOBL VARIABLES IS AVAILABLE IN THE TUTORIAL PANELS.

PLEASE SPECIFY THE FDR DIALOG CLIST LIBRARY BELOW AND PRESS ENTER:
DATA SET NAME ==> 'IDP.DIALOG.CLIST'
```

**CLIST DATA
SET NAME**

Specify the data set name of the FDR CLIST library that was loaded during installation or the library that currently contains the FDR CLISTs for the new release. This dialog will update a CLIST called ABRGLOBL which will set options every time the FDR dialogs are entered.

Note: although the dialog customization options allow for a limited ability to restrict certain dialog functions to certain TSO users or groups of users, there is another way to grant different privileges and options to varying users. If your TSO logon procedures allow the assignment of varying sets of ISPF libraries for each user, it is possible to create multiple copies of the ABRGLOBL CLIST in multiple CLIST libraries, each copy differently customized, and then allocate the proper CLIST library in the ISPF concatenations of the appropriate TSO users. Be sure that the library is one to which the users have only READ authority, otherwise they will be able to change their own FDR ISPF options and privileges; your security system may allow you to define it as "execute-only". You must also insure that they cannot provide their own copy of the ABRGLOBL CLIST in a library concatenated before the common library, for the same reason.

**DIALOG
OPTION
PANEL**

Press ENTER to display the panel containing the dialog options.

```
----- FDR INSTALLATION -- SET FDR DIALOG GLOBAL OPTIONS -----
COMMAND ==>

SAVE      - SAVE CHANGES  COPY - COPY OPTIONS          CANCEL - EXIT IMMEDIATELY

BACKGRND.. YES
FOREGRND.. SYS,OPER
REMOTEQ... YES
SCRATCH... NO
DISKUPDT.. DEFAULT      SORT..... YES          BUNIT..... TAPE

UNIT.....              SORTMSG... SYSOUT        BVOLPFX... FDRTP
XECUT..... CALL        SORTVOL... NONSPEC        RPTEXCL... $$

FSUFFIX... SPFCLIST
FVOL..... WRKPK1,WRKPK2

FDRLIB.... IDP.FDREPORT.CNTL

SORTLIB... SYS1.SORTLIB
```

90.44 CONTINUED

BACKGRND These 4 options control the execution options that users will be given for various functions generated by the FDR ISPF dialogs. If a given function has a choice of execution options, and more than one option is enabled for the user using the dialog, they will be asked to choose among them. If only one choice is enabled, it is automatically selected.

FOREGRND

REMOTEQ

SCRATCH

The options available for each are:

YES - enabled for all users

NO - disabled for all users

prefix1, prefix2, ... - a list of TSO userid prefixes which are enabled (may also include full userids)

BACKGRND and FOREGRND control background (batch job submission) and foreground (immediate TSO execution) for backup, restore, and COMPAKTOR execution. BACKGRND defaults to YES and FOREGRND defaults to a list of userid prefixes as shown above. You may want to change FOREGRND to NO.

REMOTEQ controls the ability to add ABR backup and restore requests to the remote queues. and to list and update those queues. It defaults to YES, but you should change it to NO if you have not implemented remote queue processing ([see Section 90.45](#)).

SCRATCH controls the ability to execute ABR SUPERSCRATCH from the panels. It defaults to NO. If you change it to YES or a userid prefix list, remember that SUPERSCRATCH can only be executed on volumes which are enabled for SUPERSCRATCH ([see Section 90.42](#)).

DISKUPDT Specifies the processing option for the ARCHIVE and BACKUP requests directed to the remote queues. It controls whether such requests will directly update the DSCBs of the data sets involved, or will be added to a remote queue data set (see DISKUPDATE in [Section 90.15](#)). It is used only if the REMOTEQ option authorizes this user for remote queue operations.

YES -- update the DSCB directly (requires APF authorization, [see Section 90.40](#))

NO -- Add the request to a remote queue data set.

DEFAULT -- Default to the value set in the FDR Global Option Table for DISKUPDATE ([Section 90.15](#)).

UNIT Defines the unit name to use for all allocations of disk work files created by the dialog (such as report output files). Set it to blanks (the default) to use the default unit name defined for this user (in the UADS or RACF).

XECUT Specifies the type of execution service to use when invoking a processing program in the foreground (foreground requests, reports and utilities). The options are:

CALL -- Standard CALL support. Programs will be obtained from the FDR PROGRAM library name specified in the DSNAMES Install Dialog (Option A.I.1A). This is the default.

'TSOEXEC CALL' -- TSO/E authorized CALL support (TSO/E V1.2 and above). Programs will be obtained from the FDR PROGRAM library name specified in the DSNAMES Install Dialog (Option A.I.1A).

EXEC -- ISPEXEC 'SELECT' service. Programs will be obtained from the LINKLIST, STEPLIB, or ISPLLIB data sets.

Specification of any other value will result in the value being used as a TSO command, passing the program name and parameters. For example, if XECUT is SYSCALL, the TSO command

```
SYSCALL program 'parms'
```

will be used for foreground execution of programs.

90.44 CONTINUED

- SORT** Sorting may be required for some reports generated by the ISPF dialogs. This option controls whether DD statements that may be required by your installation's SORT program should be dynamically allocated when sorting is invoked. Set SORT to **YES** to allocate the files specified by the options SORTLIB, SORTMSG, and SORTVOL below. Set SORT to **NO** if your SORT program will execute successfully with no special DD statements (many now do). The default is YES.
- SORTMSG** Defines the DDNAME to allocate for SORT output messages if SORT is YES. Specify OMIT if you do not wish to have a sort message DD statement allocated. The default is SYSOUT.
- SORTVOL** Defines the volume serial number to be used for allocation of the SORTWKnn temporary data sets if SORT is YES. Specify NONSPEC if you want the SORTWKnn data sets to go to any volume available for temporary data sets allocated from TSO. Specify OMIT if you do not wish to have the SORTWKnn data sets allocated (your SORT program may automatically allocate SORT work areas when required). The default is NONSPEC.
- BUNIT** Defines the output unit name to be use for ABR DUMPs and ARCHIVEs that are submitted for background execution. This should be a unit name for tape. (Restores are not affected by the BUNIT setting). The default is TAPE.
- BVOLPFX** If the first character of this option is "1", then the ISPF restore panels will contain a field where the user can enter a UNIT value to be used in a TAPEX DD statement for the restore. If the user leaves it blank, or BVOLPFX does not start with "1", the DYN TAPE option of ABR is used to allocate a tape drive or disk data set for the restore. The default is FDRTP.
- RPTEXCL** Defines a member name prefix from one to four characters in length. On the ISPF REPORT panel (A.1), if option 6 (FDREPORT) is chosen and the user enters an FDREPORT NAME that starts with this prefix, the user will not be required to enter other information. For other member names, the user will have to enter a data set name or volume information. The default is \$\$\$. See the notes under FDRLIB above.
- FSUFFIX** Defines the suffix for the POOLDISK data set name that describes the output disk devices to be used when executing a foreground FDRABR DUMP for backup or ARCHIVE. For more information see the description of POOLDISK option in [Section 51](#) and the option FVOL below. The full name used will be "FDRABR.POOLDISK.&FSUFFIX" (the first part may be changed by the POOLDISK option in the FDR Global Option Table, Panel A.I.5). The default is SPFCLIST.
- FVOL** Defines a list of disk volume serial numbers that will be used for the FDRABR.POOLDISK data set, which is a dummy data set that points to the volumes that are eligible to be used as output when executing foreground ARCHIVE or backup requests. If FOREGRND is enabled for this user, the ISPF dialog will check if the POOLDISK data set (see FSUFFIX above) is cataloged. If so, the volumes pointed to by that catalog entry will be used as output volumes. If not, the volumes defined by FVOL will be used. The default is 'WRKPK1,WRKPK2' but if FOREGRND is enabled and the POOLDISK name is not cataloged, this must be set to real volume serials.
- FDRLIB** Defines the name of the FDREPORT control statement data set. It is a sequential or partitioned data set containing FDREPORT control statements for predefined report formats for user execution via panel A.1 (reports). However, the SRS ISPF dialog (Panel A.S) is a far superior way of letting users execute predefined or user-specified reports, so you may want to skip this (a data set must be named, but if it does not exist, the function will be disabled). If you intend to use FDRLIB, you must create this library or name an existing library; it must have characteristics RECFM=FB and LRECL=80.
- SORTLIB** Defines the name of the SORT load module library to be allocated as DD SORTLIB if SORT is YES. Specify OMIT if your SORT program does not require a SORTLIB DD statement. The default is SYS1.SORTLIB.

90.45 CREATE THE REMOTE QUEUE FILES

NEW INSTALLATIONS: If you intend to use the ABR Remote Queues for ABR operations, you must allocate, catalog and initialize the Remote Queue data sets.

EXISTING ABR INSTALLATIONS: You will use your existing Remote Queue data sets; skip to the next section. However, you can use these instructions if you are just starting to use the Remote Queues.

THE REMOTE QUEUE DATA SETS

There are 4 possible Remote Queue data sets. Which of them you need to allocate depends on what functions you intend to let end users request via the Remote Queues, and also on the setting of the DISKUPDATE option (see [Sections 90.15](#) and [90.44](#)). Note that DISKUPDATE=YES will require that TSO be configured to allow users to execute the ABR Remote Queue utility (FDRABRUT) as an authorized program (see [Section 90.40](#)) when requesting a remote queue backup or archive; if this is not possible or desirable, you must use the data sets as shown below.

The data set names of the Remote Queue data sets are contained in the FDR Global Option Table and are set on Panel A.I.4.5 ([Section 90.16](#)). The data sets are used for:

ARCHIVE DUMP Remote Queue requests, if the DISKUPDATE option is set to NO (if DISKUPDATE is YES, this data set is not used). One index level of the name (other than the first) must be 'ABRARDQ' unless the DSNCK option (on Panel A.I.4.5) has been set to NO.

ARCHIVE RESTORE Remote Queue requests. One index level of the name (other than the first) must be 'ABRARCH' unless the DSNCK option has been set to NO.

BACKUP DUMP Remote Queue requests, if the DISKUPDATE option is set to NO (if DISKUPDATE is YES, this data set is not used). One index level of the name (other than the first) must be 'ABRBKDQ' unless the DSNCK option has been set to NO.

BACKUP RESTORE Remote Queue requests. One index level of the name (other than the first) must be 'ABRREST' unless the DSNCK option has been set to NO.

USE OF THE REMOTE QUEUE DATA SETS

The ISPF ABR dialogs (and direct execution of the Remote Queue Utility, FDRABRUT) will dynamically allocate the proper Remote Queue data set for each request and add an ABR control statement to the end of the data set (the dialogs and utility can also be used to view, change, and delete requests from the queues).

But the queues are meaningless unless you have created ABR jobs which are run by the Data Center (or Operations, or whatever central group is appropriate) which process these remote queues. The queues will be processed by ABR jobs if an appropriate DD name is added to an ABR step (see the JCL descriptions in [Sections 50](#) and [51](#)).

The DD statements for the ARCHIVE DUMP and BACKUP DUMP Remote Queues can simply be added to your regular ARCHIVE and Incremental Backup ABR jobs, to automatically include the data sets requested by the end users. However, you must be sure that the jobs will process the volumes on which the requested data sets reside (perhaps by use of the ONLVOL option of ABR).

You will probably have to create special ABR restore jobs to process the 2 RESTORE Remote Queues. The ARCHIVE RESTORE queue is the mostly likely to be used, and users will probably expect their data sets to be restored within a few hours (a day at the most) so you will need to execute the ARCHIVE Restore job at regular intervals.

ALLOCATING THE REMOTE QUEUE DATA SETS

The Remote Queue data sets should be allocated with a few tracks primary space and one track secondary with the following characteristics: RECFM=FB, LRECL=80, BLKSIZE=80, DSORG=PS. Sample member REMOTEDS in the ICL library contains JCL to allocate and initialize the remote queue data sets.

If you allocate them manually, initialize each of the remote queue data sets with an end-of-file record, e.g. by issuing the following

TSO command: `REPRO IDS('NULLFILE') ODS('remote-queue-dsname')`

90.46 INSTALLING THE FDR DIALOGS

NEW INSTALLATIONS: Although you can continue to use the ABRALLOC function shown in [Section 90.10](#) to invoke the FDR dialogs during testing, you may want to install them for use by any user while you complete your testing, to simplify testing procedures and involve a larger group of users in testing. If not, you can perform this step after testing is complete.

EXISTING INSTALLATIONS: If you have a previous version of the FDR dialogs in production, you will want to continue to use the ABRALLOC function shown in [Section 90.10](#) to invoke the new dialogs during testing. Use the procedures below to install the new dialogs when the new version is installed for production.

You may invoke the FDR ISPF dialogs by adding an option on any ISPF menu of your choosing, and/or by entering an ISPF command on any ISPF panel. For the menu option, you can either have the FDR dialog libraries allocated at TSO LOGON time, or have them dynamically allocated on demand.

ADDING THE FDR OPTION TO AN ISPF MENU

You may add FDR as an option on the ISPF main menu or any menu of your choice. In the FDR dialog panel data set, there are four sample panels showing how to add ABR to the ISPF main menu (ISR@PRIM):

ISR@V3X - for ISPF V3.x main menu (non-CUA) using pre-allocated libraries.

ISR@V3XL - for ISPF V3.x main menu (non-CUA) using dynamically allocated libraries.

ISR@V4X - for ISPF V4.x main menu (CUA) using pre-allocated libraries.

ISR@V4XL - for ISPF V4.x main menu (CUA) using dynamically allocated libraries.

You can make similar modifications to any other ISPF menu. If you add ABR to a non-CUA panel, follow the instructions in ISR@V3X or ISR@V3XL; if you add ABR to a CUA panel, follow the instructions in ISR@V4X or ISR@V4XL. Note, however, that the ABR version of the main tutorial panel (ISR00003) assumes that the ABR option appears on the ISPF main menu.

WARNING: Please do not code LIBDEF statements. Please use the technique shown in members ISR@V3XL and ISR@V4XL, invoking CLIST ABRALLOC. Otherwise, various errors will probably occur.

ADDING THE FDR DIALOG LIBRARIES TO THE LOGON PROCEDURE

If you choose to have the FDR Dialog libraries allocated by the TSO LOGON procedure, instead of dynamically allocated on demand, then the FDR dialog libraries should be allocated as the first library of the concatenation of the following DDNAMEs (the default data set name of each library is shown; if you have changed the names during installation, substitute your names):

SYSPROC -- TSO command procedure libraries (CLISTs), add 'IDP.DIALOG.CLIST'.

ISPLLIB -- ISPF dialog panel libraries, add 'IDP.DIALOG.PANELS'.

ISPMMLIB -- ISPF dialog message libraries, add 'IDP.DIALOG.MESSAGES'.

ISPSLIB -- ISPF dialog skeleton libraries, add 'IDP.DIALOG.SKELETON'.

ISPLLIB -- ISPF TASKLIB library, add 'IDP.MODFDR53'. If your TSO LOGON proc contains a STEPLIB DD statement but does not contain an ISPLLIB DD statement, then add the FDR load module library to the STEPLIB DD concatenation. However, there is no need to allocate the FDR load module library in the LOGON procedure if it has been placed in the system LINKLIST.

90.46 CONTINUED

**ADDING FDR
COMMANDS
AS ISPF
COMMANDS**

You can add FDR commands to an ISPF command table, so that users will be able to invoke the FDR/ABR dialogs and services by entering the command name (followed by any applicable parameters) on the command line of almost any ISPF panel. If you wish to do so, then from the FDR Installation Options Menu (usually accessed by A.I), enter 11 (ADDISPF) to display this panel:

**PANEL A.I.11
ADD FDR
COMMANDS
AS ISPF
COMMANDS**

```

----- Add FDR Commands to an ISPF Command Table  Row 1 to 20 of 20
COMMAND ==>                                     SCROLL ==> PAGE

      ADD - add the selected table entries      CANCEL - exit immediately

How are the FDR dialog libraries allocated ?
Allocation      ==> 2      (1 - at TSO LOGON time, by the LOGON proc or CLIST)
                        (2 - on demand, by the ABRALLOC CLIST)

ISPF table library to update:(optional)
Data set name   ==>
ISPF command table to update:
Application ID   ==>
Sel Command Truncate Type Description
-----
S   FDR          3      PRIM INVOKE THE FDR DIALOG
S   SRS          3      SRS  PERFORM SRS DATA SET SEARCH
S   S2           2      S2   DISPLAY SRS DATA SET LIST
   ARCDL         4      FF   DELETE DATA SET ENTRY IN THE ARCHIVE FILE
   ARCHIVE       3      FF   ARCHIVE DATA SET
   ARCMOD        4      FF   MODIFY DATA SET ENTRY IN THE ARCHIVE FILE
   ARCRECAT      6      FF   RECATALOG ARCHIVED DATA SET FOR AUTO-RECALL
   ARCRESET      6      FF   RESET DATA SET ENTRY IN THE ARCHIVE FILE
   BACKAPPL     5      FF   DATA SET APPLICATION BACKUP
   BACKUP        3      FF   ADD BACKUP REQUEST TO REMOTE QUEUE
   COPYDS       5      FF   COPY DATA SETS
   MOVEDS        5      FF   MOVE DATA SETS
   REMOVEA      7      FF   DELETE ARCHIVE RESTORE REQUEST FROM REMOTE QUEUE
   REMOVEB      7      FF   DELETE BACKUP RESTORE REQUEST FROM REMOTE QUEUE
   REORG        3      FF   COMPRESS PDS DATA SETS
   RESETARC     6      FF   DELETE ARCHIVE REQUEST FROM REMOTE QUEUE
   RESETBKP     6      FF   DELETE BACKUP REQUEST FROM REMOTE QUEUE
   RESTAPPL     6      FF   RESTORE DATA SETS FROM APPLICATION BACKUP
   RESTARC      5      FF   RESTORE DATA SETS FROM ARCHIVE
   RESTBKP      5      FF   RESTORE DATA SETS FROM BACKUP
***** Bottom of data *****

```

How are the FDR dialog libraries allocated ? ==> specifies whether the FDR dialog libraries (1) are allocated at TSO LOGON time, by the LOGON proc or LOGON CLIST, or (2) are allocated on demand, by the ABRALLOC CLIST.

ISPF table library to update: specifies the name of a table library, allocated to the ISPTLIB ddname concatenation, where the command table will be updated. The default is that the specified command table will be updated in the first library in which it exists in the ISPTLIB concatenation for this TSO userid.

ISPF command table to update: Application ID ==> specifies the ISPF command table name prefix (omit the command table name suffix, CMDS). The default is ISP, for the ISPF system command table ISPCMDS. For ISPF version 4.2 or higher, it will be preferable in most cases to specify the site command table, so that the FDR commands will not need to be re-installed for new releases of ISPF. The installation specifies the prefix for the site command table in the ISPF configuration table ISRCONFIG, as documented in the IBM manual ISPF Planning and Customizing or equivalent. The Help for panel A.I.11 shows the prefix for the site command table at your installation, and whether the site command table is searched before or after ISPCMDS.

Sel specifies whether the entry is selected to be added: S or blank.

Command specifies the command name to be added to the command table.

90.46 CONTINUED

Truncate optionally specifies the number of characters for abbreviation; the minimum value is 2.

Type specifies one of the following:

PRIM if the command invokes the FDR Primary options panel;

SRS if the command invokes an SRS Data Set Search;

S2 if the command invokes an SRS Saved Data Set List;

FF if the command invokes an SRS Function command.

Description command description.

Note that only the following commands will be of use to non-ABR customers: FDR, COPYDs, MOVEDs, and REOrg; also that the REOrg command will only be of use to customers with the FDRREORG component.

The FDR ISPF commands (except for FDR, SRS, and S2) invoke FDRSRS function commands. Each SRS function command contains a set of pre-defined parameters, which are fully customizable. The user may also create new SRS function commands as described in [Section 54](#).

**RESTRICTING
THE USE OF
THE INSTALLA-
TION PANELS**

Since you probably want only certain users to be able to execute the FDR Installation dialogs, you can edit member ABRINPRI in the FDR Dialog CLIST library. ABRINPRI contains instructions for restricting user access.

To insure that users don't copy and modify ABRINPRI or other FDR CLISTs, you may be able to use your security system to grant execute-only access to the FDR Dialog CLIST library for unauthorized users.

**CHANGING
THE DATA SET
NAMES OF THE
FDR LIBRARIES**

The names of the FDR libraries that are used in the ISPF dialogs may be changed by using the DSNAMES dialog. From the FDR Installation Options Menu (usually accessed by A.I), enter 1A (DSNAMES) to display this panel:

**PANEL A.I.1A
CHANGE
DSNAMES
USED BY
DIALOGS**

```

----- FDR INSTALLATION -- CHANGE DATA SET NAMES -----
COMMAND ===>

      SAVE      - SAVE DATA SET NAMES      CANCEL - EXIT IMMEDIATELY

PLEASE SPECIFY THE DATA SETS NAMES OF THE FDR LIBRARIES. PRESS ENTER TO VERIFY.

CMD LIBRARY  DATA SET NAME (QUALIFIED)                                     ( CMD: Browse Edit View )
-----
CONTROL      IDP.ICLFDR53
PROGRAM      IDP.MODFDR53
CLIST        IDP.DIALOG.CLIST                                           (required)
PANELS       IDP.DIALOG.PANELS
MESSAGES     IDP.DIALOG.MESSAGES
SKELETON     IDP.DIALOG.SKELETON
JCL          IDP.JCLFDR53
-----

STEPLIB to the FDR Program Library ===>      YES (yes no) - use no if in LINKLIST

Note: The values displayed above originated from the current CLIST allocations,
      DSNAMES=clist library name
    
```

If the DSNAMES dialog is not available, invoke the CLIST directly, by entering the TSO command

```
EXEC 'IDP.DIALOG.CLIST(ABRALLOC)' 'DSNAMES'
```

in TSO READY mode or in ISPF option 6.

Note: Other FDRABR data set names are specified on the ABR Data Set Names panel A.I.4.5 ([section 90.16](#)).

90.47 LOAD SAR ONTO DISK

NEW INSTALLATIONS: Most installations will want to place a copy of SAR onto a disk volume for rapid recovery from system outages due to DASD failures.

EXISTING INSTALLATIONS: You can use this dialog to place SAR onto new disk volumes, or to create a copy of SAR on tape.

SAR (Stand-Alone Restore) is the component of FDR which can be IPLed without an operating system, to restore backups, create backups, or change the volume serial of disk volumes. SAR and FDRSARLR are documented in [Section 15](#).

This dialog generates a batch jobstream for the SAR Loader program (FDRSARLR) which will load SAR onto a disk volume from which it can be IPLed, and optionally change the SAR default responses to the values you use most often. On disk, SAR actually occupies no space; it is placed in the "IPL text" area of the label track (Cylinder 0 Track 0) on the volume; you load SAR by IPLing from the disk volume. This dialog can also be used to copy SAR to a tape (by editing the generated jobstream to change the output).

Innovation strongly recommends SAR be loaded onto one or more disk volumes, preferably several volumes. Be sure to document the SAR disk addresses and volume serial numbers somewhere convenient in the computer room.

PANEL A.I.2 SAR LOADER UTILITY

Select option 2 (LOADSAR) on the FDR Installation Options Menu which will bring you to this panel

```

----- SAR LOADER UTILITY -- ENVIRONMENTAL SPECIFICATIONS -----
COMMAND ==>

SEL   - SELECT SAR DEFAULTS PANELHELP   - TUTORIAL HELP PANELS

OUTPUT DISK DEVICE: (3380/3390)
VOLUME SERIAL      ==>                      REWRITE ==> NO   (YES/NO)

INSTALLATION CONTROL LIBRARY DATA SET:
DATA SET NAME      ==> 'IDP.ICLFDR53'

FDR PROGRAM LIBRARY DATA SET:
DATA SET NAME      ==> 'IDP.MODFDR53'

SYSOUT CLASS       ==> *

JOB STATEMENT INFORMATION:
==> //useridA JOB (ACCOUNT),'NAME'
==> // *
==> // *

```

When you have completed the fields shown below, enter SEL on the command line to display another panel where you can SUBMIT or EDIT the batch FDRSARLR job and optionally set defaults for SAR responses.

OUTPUT DISK DEVICE

Specify the volume serial of the disk volume where SAR will be written. If it already contains an earlier copy of SAR, specify REWRITE as YES to allow FDRSARLR to overwrite the earlier version.

LIBRARY DATA SET NAMES

If the name of the FDR Installation Control Library (ICL) or program library is not correct, correct them. The ICL contains a copy of SAR used as input to FDRSARLR.

JCL INFORMATION

Since a batch jobstream will be generated to execute FDRSARLR, specify SYSOUT CLASS and JOB STATEMENT information appropriate for running the job.

90.50 TESTING THE FDR SYSTEM

This section provides guidelines for testing various components of the FDR system. Use these guidelines to create and execute a test plan for testing the new version of FDR, testing the FDR components which you have installed. Customers trialing FDR will have all components of FDR installed; we hope that you will try them all, but, of course, you can limit testing to only those components you intend to license.

Customers with ABR can test, at this point, everything except functions requiring the ABR Operating System Exits, primarily auto-recall of ARCHIVED data sets and automatic restore from backup of scratched data sets. Customers with FDRREORG can test everything except transparent intercept of IEBCOPY compress requests. Instructions for testing these functions with the new exits are in [Section 90.62](#).

All of the example jobstreams shown in the various sections of this manual may be found in the JCL library loaded as part of the FDR installation. The members names include the section number of the manual in which the example is found so that you can easily find an example that you see in the manual.

**TESTING FDR,
FDRDSF,
FDRCOPY**

NEW INSTALLATIONS: You can test the basic functions of backup, restore, and copy/move with simple batch jobstreams. If FDR has not yet been installed in the system linklist, a STEPLIB DD statement must be used to point to the FDR library.

EXISTING INSTALLATIONS: You can test the basic functions of backup, restore, and copy/move with simple batch jobstreams. A STEPLIB DD statement must be used to point to the FDR library.

FDR full-volume backups and restores can be tested using the sample jobstreams found in [Section 10](#) (FDR). You can backup any volumes you like (FDR backups make no changes to the input volume) but if you backup live data, be sure you restore to a spare scratch volume with the CPYVOLID=NO operand to preserve the output volume serial and avoid back-leveling the live volumes. Note that the CPYVOLID=NO restore creates uncataloged copies of the restored data sets; non-VSAM data sets can be accessed by specifying the volume serial of the scratch volume but VSAM clusters on the scratch volume will not be usable.

FDRDSF data set backups and restores can be tested using the sample jobstreams found in [Section 20](#) (FDRDSF). You can backup any data sets you like (DSF backups make no changes to the input data sets) but if you backup live data, be sure that you restore data sets to new names (the NEWNAME= or NEWINDEX= operands on the SELECT statement) to avoid back-leveling the live data sets. You may wish to test restoring individual data sets with FDRDSF from full-volume backup tapes created by FDR. You may want to browse the restored data sets or run other programs to verify their contents.

You may want to test backup to and restore from backup data sets on disk as well as backups on tape.

Note that FDR and FDRDSF use the same backup and restore modules that are used internally by ABR backups and restores, so by testing FDR and FDRDSF you are also testing much of the code used by ABR.

You can test FDRCOPY COPY and MOVE of data sets using sample jobstreams found in [Section 21](#) (FDRCOPY). You can copy any data sets you like (COPY makes no changes to the input data set) but be sure to specify new output names (the NEWNAME= or NEWINDEX= operands on the SELECT statement). MOVE tests should be done on data sets specially created for the test since MOVE scratches the input data set; you might use COPY to create test data sets for MOVE. You may want to browse the copied data sets or run other programs to verify their contents.

90.50 CONTINUED

**TESTING
FDRREORG**

NEW INSTALLATIONS: You can test all functions of FDRREORG except transparent intercept of IEBCOPY compress requests. If FDR was not installed in the LINKLIST, a STEPLIB DD statement should be used to point to the FDR library.

EXISTING INSTALLATIONS: You can test all functions of FDRREORG except transparent intercept of IEBCOPY compress requests. A STEPLIB DD statement should be used to point to the FDR library containing the new version.

You can test the FDRCOPY REORG function, to unconditionally compress PDSs, using the sample jobstreams found at the end of [Section 32](#). The SIMREORG statement can be used to simulate compression of any PDSs, reporting on the results that would be accomplished by a real compression. You can also execute a live REORG against test data sets. Note that an FDRCOPY COPY of a PDS will copy that PDS in its existing state, without reorganizing it, so you can copy PDSs needing compression, then test the REORG of that copy. Compare utilities (such as IEBCOMPR or ISPF option 3.12) can be used to verify that the members are unchanged after the compression.

You can test the FDRREORG functions using sample jobstreams found in [Section 30](#) (FDRREORG). The SIMULATE (SIM) statement of FDRREORG can be used to simulate reorganization, but note that SIM normally only reports on the data sets it would select for reorganization. You can specify the SIMPDSCOMP operand on the SIMULATE statement to invoke FDRCOPY to do a SIMREORG of any selected PDSs. You can also do a real REORG with the NOREORG operand; this will cause FDRREORG to actually read and create backups for selected VSAM and IAM data sets but not to read those backups back for reorganization, thus testing the backup process.

Actual reorganizations should be tested with care, selecting data sets or volumes which have backups or can be recreated, if necessary.

TESTING SAR

Many installations neglect to test SAR (Stand-Alone Restore). It is the most awkward component of FDR to test, but when SAR is required to recover from a hardware failure or disaster, it may be the most critical component.

Before scheduling a SAR test, you should create a non-labeled tape containing SAR, and also place IPLable copies of SAR on one or more disk volumes, using the instructions in [Section 90.47](#) or [15.20](#). You can IPL SAR from the FDR distribution tape, but this requires that you IPL 5 times to bypass the tape header labels. IPLing from an unlabeled tape or disk is faster and more convenient.

SAR must be tested on a dedicated CPU or LPAR, so test time will probably have to be scheduled. SAR can also be tested on a VM virtual machine, but unless this is the environment that SAR will actually be used, it is not recommended since it may not reflect your real recovery environment that SAR will face.

Follow the instructions in [Section 15](#) for loading and executing SAR. You should test the full-volume backup and restore functions of SAR. You will probably want to be sure that SAR can restore from a tape created by FDR, and that FDR can restore from a tape created by SAR (note that FDRDSF can restore non-VSAM data sets from a SAR full-volume backup, but not VSAM clusters).

Be sure to restore your test backups to spare (scratch) output volumes. During the SAR restores, you can specify the CPY=N option (equivalent to the FDR CPYVOLID=NO operand) to preserve the output volume serial during the restore.

90.50 CONTINUED

**TESTING
COMPAKTOR**

NEW INSTALLATIONS: You can test all functions of COMPAKTOR. If FDR was not installed in the LINKLIST, a STEPLIB DD statement should be used to point to the FDR library.

EXISTING INSTALLATIONS: You can test all functions COMPAKTOR. A STEPLIB DD statement should be used to point to the FDR library containing the new version.

You can test COMPAKTOR using sample jobstreams found in [Section 40](#) (COMPAKTOR). Most COMPAKTOR functions can also be tested using the A.C (COMPAKTOR) and A.R (RELEASE) options on the FDR Primary Options Menu panel.

All COMPAKTOR functions except MAP can be simulated with the SIMULATE (SIM) statement. SIMULATE will report on exactly what volume changes will be made by the equivalent real COMPAKTOR function; for a FAST COMPAKTion (FASTCPK) it will also estimate the elapsed time of the real function (based on the number of tracks which must be moved).

TESTING ABR

NEW INSTALLATIONS: You can test all functions of ABR except functions requiring the ABR Operating System Exits, primarily auto-recall of ARCHIVED data sets and automatic restore from backup of scratched data sets. However, you can create ABR backups and ARCHIVES now which can be used when testing the exits later ([Section 90.62](#)). If FDR was not installed in the LINKLIST, a STEPLIB DD statement should be used to point to the FDR library.

EXISTING INSTALLATIONS: You can test all functions of ABR except functions requiring the ABR Operating System Exits (but some functions may work with your existing exits). You must take care that the ABR functions you test do not impact your production ABR environment; this is especially important for full-volume and incremental backups since any backups you create of volumes currently being backed up by ABR will become part of the production ABR system. This is not a problem for ABR application backups. ARCHIVES can be tested using a separate ARCHIVE Control File. A STEPLIB DD statement should be used to point to the FDR library containing the new version.

**TESTING ABR
VOLUME
BACKUPS**

ABR volumes backups can be simulated (SIM statement). This will show you the volumes and data sets that will be processed. But since no actual backup is performed, this cannot be used to test the actual operation of incremental backups. A SIM of a volume which has never been actually processed by ABR volume backup will always simulate a full-volume backup.

Full-volume and incremental backups can be tested using the sample jobstreams found in [Section 50](#) (ABR VOLUME BACKUPS). You will probably want to setup a few special test volumes for these tests. Initialize the volumes for ABR processing using Panel A.I.8 ([Section 90.42](#)); be sure to enable OLDBACKUP.

The first backup of each test volume must be an ABR full-volume backup (DUMP TYPE=FDR). Actually, the full-volume backup is forced if you try to do an incremental backup. Incremental backups (TYPE=ABR or TYPE=AUTO) can be executed after that. Also, incremental backups are usually executed once a day, you can execute as many as you like on any day to simulate the normal operation of ABR.

You will probably want to update a number of data sets on the volumes between each backup; otherwise incremental backup will have nothing to do. You can do this with ISPF EDIT or any other program which opens files for output.

You may want to execute enough cycles of full and incremental backups to verify that the old ABR generations are uncataloged from the ABR cataloged (based on the GEN value used to initialize the volume).

You will want to test data set restores and full-volume recoveries. Data set restores can be executed in batch (using sample jobstreams from [Section 50](#)), but they can also be executed from the FDR ISPF dialogs, using either the ABR RESTORE panel (A.2) or the SRS dialog (A.S) which allows you to display the backups available for data sets and restore any of them.

90.50 CONTINUED

**TESTING ABR
APPLICATION
BACKUPS**

ABR application backup (TYPE=APPL) is easy to test, since the backups can be run against any data sets or volumes. Application backup has no effect on the data sets backed up so you can test against live data. There are no special initialization requirements. The sample jobstreams in [Section 52](#) (ABR APPLICATION BACKUPS) can be used to setup tests.

When restoring from application backup tests, you should probably restore to new names (using the NEWINDEX= operand on the SELECT statement) and may want to direct the restored data sets to new volumes (the NVOL= operand) to avoid overlaying any live data.

**TESTING ABR
ARCHIVE AND
SUPER-
SCRATCH**

ABR ARCHIVE (TYPE=ARC) and SUPERSCRATCH (TYPE=SCR) should be tested against data sets and volumes created specially for the test, since both functions will scratch the data sets from the disk volumes. The jobstreams in [Section 51](#) (ABR ARCHIVE) can be used to setup tests.

For existing customers with ARCHIVING already in production, you should setup an ARCHIVE Control File just for testing, to avoid mingling test records with live records in the normal Control File. The test Control File can be created using ISPF Panel A.I.9 ([Section 90.43](#)); the name of the Control File should be set in the FDR Global Options Table in the test FDR program library (Panel A.I.4.5, [Section 90.16](#)) so that the DYNARC operand of ABR can be used. This Control File can also be used for auto-recall testing when testing the ABR exits as in [Section 90.62](#).

You can test SUPERSCRATCH by creating some test data sets (VSAM and non-VSAM) and then selecting them in a DUMP TYPE=SCR test. The data sets should be deleted and uncataloged.

You can test ARCHIVE by creating some test data sets (VSAM and non-VSAM) and then selecting them in a DUMP TYPE=ARC test. The data sets should be backed up, recorded in the ARCHIVE Control File, deleted and uncataloged (except that if you run with the RECALL=YES option specified or defaulted, they will remain cataloged with a special auto-recall indicator; if you also specify or default to MIGRAT=YES, the volume serial in the catalog is changed to MIGRAT).

You should test ARCHIVE with the backups directed to tape, and also to disk.

You can then test restoring some of the ARCHIVED data sets, both singly and collectively (several together), restoring from disk and tape backups. You can submit batch jobstreams to restore (using examples from [Section 51](#)) or you can use the FDR ISPF dialogs using either the ABR RESTORE panel (A.2) or the SRS dialog (A.S) which allows you to display the ARCHIVES available for data sets and restore any of them.

You should also test the FDRARCH REORG function ([Section 51.40](#)) to reorganize the ARCHIVE Control File. You can use the FDRARCH MODIFY statement to change the recorded expiration dates of a few of your ARCHIVED data sets (or the DELETE statement to flag them for deletion), so that you can see the REORG operation deleting the records and reclaiming the space.

**TESTING
FDREPORT**

You may want to test FDREPORT, the generalized reporting program. FDREPORT makes no changes to any input volume, so you can run any reports you desire. The examples in [Section 54](#) (FDREPORT) can be used, but you may also want to use the Innovation Health Check jobstreams in the FDR JCL Library, member names HCHECKx, both to test FDREPORT and to give you valuable information about your DASD system (see [Section 54](#) for more information).

The SRS (Search, Report, Services) ISPF dialog (panel A.S) uses FDREPORT internally for most of its information gathering. You can execute a variety of SRS displays to test both SRS and FDREPORT. SRS is described in [Section 54](#).

90.60 DYNAMIC INSTALLATION OF THE FDR EXITS

NEW INSTALLATIONS: If your FDR system includes ABR and/or FDRREORG, you will probably want to activate the appropriate FDR system exits, using the dynamic exit installation described in the following sections. These exits can be installed and deactivated as required during your testing. When testing is complete, follow the procedures described to insure that the exits are installed following every IPL.

EXISTING INSTALLATIONS: If your FDR system includes ABR and/or FDRREORG, you will want to test the new version of the exits. Procedures for testing the new exits while continuing to run the exits from the prior release of FDR are described in the following sections; you can restrict the test exits to certain users or jobs. Once the new FDR is installed for production, be sure to follow the procedures for activating the new exits for use by all users and be sure that the new exits are installed after every IPL.

The FDR system includes several Operating System exits. Although these exits are not required for ABR to function, they provide very important additional functions. They are installed by a dynamic installation process which is transparent to other system functions. The FDR/ABR Operating System exits are installed at well-defined exit and interface points provided by MVS whenever possible, and are designed to co-exist with other exits which may already be in use at the same exit points.

The dynamic installation process enables the FDR Operating System exits to be automatically installed at IPL time and to remain throughout the life of the IPL. Once installed, the exits may be deactivated and reactivated, possibly with different installation options.

**EXITS THAT ARE
INSTALLED
DYNAMICALLY**

The FDR/ABR MVS exits that are available for dynamic installation are the CATALOG LOCATE exit and the Data Set Not Found (DSNF) exit, which implement the Automatic Recall feature of ABR, the DADSM Pre-processing exit and the FDRREORG IEBCOPY intercept.

**EXITS FOR
AUTOMATIC
RECALL**

The ABR CATALOG LOCATE exit (module FDR026DU) and the ABR Data Set Not Found exit (module FDREXDSN) implement the Automatic Recall feature of ABR, which is discussed in detail in [Section 51](#). With these exits installed, a reference to a data set that has been ARCHIVED, from either a batch job or a TSO user, will cause the data set to be automatically restored. If these exits are not installed, any restore of an ARCHIVED data set must be explicitly requested.

When the LOCATE exit is installed, two other supporting exits are also installed: FDRDAS02 (a DADSM scratch exit to suppress irrelevant messages associated with deleting ARCHIVED GDGs) and FDR00024 (to implement transparent support for recall requests using the HSM interfaces). Also, if there is no module called ARCGIVER (a HSM module involved in the HSM interface) a module called FDRGIVER is installed.

90.60 CONTINUED

**DADSM PRE-
PROCESSING
EXIT**

The purpose of the ABR DADSM Pre-processing exit (module FDRPRE00) is to record ABR backup information in the ABR SCRATCH catalog for any data sets that are scratched (deleted) or renamed by programs other than FDRABR if they have current ABR backups. With the DADSM Pre-processing exit installed, a user can request the restore of a data set that has been scratched or renamed simply by specifying the data set name. If this exit is not installed, the backups are still available, but the user must tell ABR the generation and cycle from which to restore a scratched data set.

When a data set is renamed, the DADSM Pre-processing exit clears the ABR backup indicators to indicate that the data set has not been backed up under the new name. When a new data set is allocated, or a data set is extended to a new volume, this exit initializes the last reference date.

When FDRPRE00 is installed, a second exit (FDRPOST0, a DADSM Post-Processing exit) also appears to be installed; actually it is an additional entry point in FDRPRE00.

Innovation strongly recommends installation of the DADSM Preprocessing exit.

**FDRREORG
IEBCOPY
INTERCEPT**

FDRREORG can provide an optimized replacement for an IEBCOPY compress-in-place operation providing substantial savings in elapsed time and resource consumption. This intercept enables FDRREORG to transparently replace IEBCOPY compress operations. When installed, any invocation of IEBCOPY will actually invoke the intercept; the intercept will determine if a compress-in-place was requested by the user; if so, an FDRCOPY REORG is invoked to compress the data set. Other IEBCOPY functions are passed to the real IEBCOPY. The intercept will also handle programs which call IEBCOPY for compression, such as SMP/E, and ISPF.

90.61 DYNAMIC EXIT INSTALLATION PROCEDURE

The following section describes the procedure for the dynamic installation of the ABR and FDRREORG exits.

For existing installations testing a new version while an existing version of FDR and the exits is still in use, please review [Section 90.62](#) for procedures used to test the new exits in parallel with production.

**STEP1: SET
EXIT OPTIONS**

If not already done, set the dynamic installation options via the ISPF panel A.I.4.11.1 ([Section 90.22](#)). On this panel there are separate options to enable installation of each of the ABR and FDRREORG exits. **They are all disabled by default; at least one must be enabled for the dynamic installation to function. Innovation recommends that ABR customers enable all the ABR exits.**

The panel also allows you to choose whether the system operator must confirm, by replying to a console message, that the exits are to be installed. You may wish to enable this during testing, but in production it requires a reply as part of every IPL which may not be desirable.

Various options may also be set to modify the processing done by the ABR exits. These options are set on ISPF panels:

AUTO-RECALL (mostly the LOCATE exit) options on Panel A.I.4.11.2 and A.I.4.11.3

DATA SET NOT FOUND (DSNF) exit options on Panel A.I.4.11.4 (currently has no options)

DADSM Preprocessing exit options on Panel A.I.4.11.5

For details, [see Section 90.22](#).

If you change any of the options that affect the exits, after the exits are dynamically installed, and you want the changes to take effect before the next IPL, you must run FDRSTART, as discussed in Step 7.

**STEP 2:
LINKLIST**

Use of FDR with the exits in production requires that the FDR load module library is in a system LINKLIST library that is APF authorized. You must either add the FDR program library to the linklist, or copy the FDR modules to an authorized linklist library.

**STEP 3:
SYNRECAL
PROCEDURE**

Copy the cataloged procedure SYNRECAL from the FDR Installation Control Library (ICL) to a cataloged procedure library that is available to JES for START commands. Edit this member and remove the STEPLIB DD statement, since the FDR load module library is in the LINKLIST. You may rename this cataloged procedure to a name other than SYNRECAL; if so, change the value for the LXSYPNPROC option on ISPF panel A.I.4.11.2.

90.61 CONTINUED

STEP 4: During initial testing, the exits may be dynamically installed by executing a batch job using the following JCL:

ACTIVATE EXITS

```
//FDRSTART    JOB          'DYNAMICALLY INSTALL FDR EXITS',FDR
//FDRSTART    EXEC        PGM=FDRSTART
//ABNLDUMP    DD          DUMMY                      FOR ABEND-AID
//SYSUDUMP    DD          SYSOUT=A
```

This JCL is supplied in member FDRSTART in the FDR Installation Control Library (ICL). You can also generate this job using the REFRESH command on ISPF Panel A.I.4 (see [Section 90.11](#)).

Note: In case of abnormal termination or errors, the dynamic installation program error recovery routine will issue a diagnostic message to the console operator and deactivate any exits which were installed.

The exits (and some supporting code) will be loaded in common memory (CSA), and various system pointers will be updated to activate them.

During testing, run jobs to test the function of the FDR/ABR exits. Please consult [Section 90.62](#) for test guidelines.

After running the FDRSTART job, the exits will remain active until the next IPL of your Operating System. After an IPL, you must run the job again to reactivate the exits for testing. When you are ready to have the exits activated automatically during each IPL, see Step 7.

The status of the dynamically installed FDR/ABR exit modules may be obtained by executing a batch job using the following JCL:

```
//FDRSTATS    JOB          'REPORT STATUS OF FDR EXITS',FDR
//FDRSTATS    EXEC        PGM=FDRSTATS
//ABNLDUMP    DD          DUMMY                      FOR ABEND-AID
//SYSUDUMP    DD          SYSOUT=A
```

STEP 5: This step will deactivate (but not deinstall) the exits; it is not normally needed but can be used during testing if desired.

DEACTIVATE EXITS

If the dynamically installed exits are deactivated, their processing will be completely bypassed, causing all requests to be ignored, as if the exits were not installed. An IPL is required to completely remove the exits from the system. Normally, the exits should remain active for the life of the IPL. There is no need to deactivate the exits when shutting down the system.

All of the dynamically installed exit modules may be deactivated by executing a batch job using the following JCL:

```
//FDRPARE     JOB          'DEACTIVATE FDR EXITS',FDR
//FDRPARE     EXEC        PGM=FDRPARE
//ABNLDUMP    DD          DUMMY                      FOR ABEND-AID
//SYSUDUMP    DD          SYSOUT=A
```

This JCL is supplied in member FDRPARE in the FDR Installation Control Library (ICL).

The exit modules may be reactivated by re-executing the FDRSTART program.

Individual exits may be deactivated without interfering with the operation of other exits that you want to keep running, by turning off the appropriate dynamic installation option (ABRLOC, ABRDSNF, ABRPRE and/or IEBCOPY) on panel A.I.4.11.1 and rerunning FDRSTART. An exit that has been deactivated in this way may be reactivated by changing the dynamic installation option back and rerunning FDRSTART again.

90.61 CONTINUED

**STEP 6:
PRODUCTION**

After preliminary testing, the dynamic installation of the FDR/ABR exits should automatically take place as part of each IPL. You could do this by executing the FDRSTART job automatically during the IPL procedure, but the install will be more efficient (and done earlier in the IPL process) if you make FDR a subsystem.

To accomplish this, add a record in one of these formats to an IEFSSNxx member in data set SYS1.PARMLIB:

1) Positional format (starting in column 1):

```
FDR , FDRSTART
```

2) Keyword format (supported in MVS/ESA 5.2 and OS/390):

```
SUBSYS SUBNAME(FDR) INITRTN(FDRSTART)
```

You may place this record into a new IEFSSNxx member and add the new member's suffix to the SSN= parameter in IEASYSxx, or you may add it to an existing IEFSSNxx member (but you must use the same format as the other records in that member). This will cause the FDRSTART program to run during every IPL, as a subsystem initialization routine, to dynamically install the FDR Operating System exits. In the current release, FDR will not actually be started as an MVS subsystem; it will terminate after the exits are installed.

For the first IPL after updating IEFSSNxx, it is a good idea to set the CONFINST option on panel A.I.4.11.1 to cause the console operator to be asked to confirm that the dynamic installation of the ABR Operating System exits should take place. Then, in the unlikely case that the exits cause the IPL to fail, you can re-IPL and bypass the exit installation. Once you are satisfied that the exits perform properly, you can disable CONFINST.

**STEP 7:
CHANGING
OPTIONS**

FDRSTART will also load a copy of FDROPT, the FDR Global Option Table, into common memory (CSA) with the exits. The FDROPT module will appear to be in LPA so options are used by the exits, as well as by any executions of FDR programs from the system linklist (without a STEPLIB). If you change any of the options in the Option Table after the FDR Operating System exits are dynamically installed, and you want the changes to take effect before the next IPL, follow this procedure to refresh the in-memory copy of FDROPT.

A. Change the desired options, using the ISPF Installation Panels (A.I.4.x), updating the copy of FDROPT in the FDR program library in the system linklist.

B. Execute FDRSTART in a batch job, using the JCL shown in Step 4 or using the REFRESH command on installation panel A.I.4. This may be done at any time, even if the exits are normally installed at IPL time using IEFSSNxx as shown in Step 6. When program FDRSTART sees that the dynamically installed exits have already been activated, it does not reload the exit modules, but it does update FDROPT. If you have more than one system, do this on every system.

NOTES:

- 1) This procedure is not required if the only exit installed is the FDRREORG IEBCOPY Intercept.
- 2) In releases prior to V5.3, it was necessary to do an LLA REFRESH after updating the options before running FDRSTART. This is no longer necessary; FDRSTART will refresh the LLA copy of FDROPT before loading it. FDRSTART will also refresh the LLA copy of all FDR tables and lists, including the COMPAKTOR Unmovable Table, the ABR Restore Allocate List, the ABR Protect Lists, and the FDRREORG NOREORG List, in case they have been updated.
- 3) In the FDR ISPF Dialogs, Option A.I.4A takes you to a version of the install dialog where the options currently set in the in-memory copy of FDROPT can be displayed but not modified. You can use this to verify that the refresh procedure has worked successfully.

90.61 CONTINUED

**STEP 8:
UPDATING
THE EXITS**

If you need to update the installed exits before the next IPL, replacing them with new versions (perhaps to apply a fix from Innovation), you can do so with a batch job using the following JCL:

```
//FDRSTART    JOB          'DYNAMICALLY REPLACE FDR EXITS',FDR
//FDRSTART    EXEC        PGM=FDRSTART,PARM='PROD,REPLACE'
//ABNLDUMP    DD          DUMMY                                FOR ABEND-AID
//SYSUDUMP    DD          SYSOUT=A
```

This JCL assumes that the new version of the exit modules resides in LINKLIST. You can also generate this job with the REFRESH command on Panel A.I.4 (the REPLACE parm can be specified on the FDRSTART panel).

This may be done at any time, even if the exits are normally installed at IPL time using IEFSSNxx as shown in Step 6. FDRSTART will reload a new copy of the exit modules and options table, replacing the exits and options table that were previously installed. This job should be run on every system on which you want to install the new version of the exits.

Note: The REPLACE keyword causes FDRSTART to allocate CSA storage for the new set of exits. This is in addition to the CSA storage already occupied by the replaced exits, which remains allocated. Therefore, REPLACE must not be used indiscriminately in order to prevent CSA shortage problems (requiring a re-IPL). If Innovation supplies you with a zap which affects the exit modules, you will also be given a procedure which can be used to apply that zap to the installed exits in-place without requiring an IPL or using the REPLACE procedure.

90.62 TESTING THE FDR MVS EXITS

Here are some suggestions for tests that you may run to verify that the FDR Operating System exits are installed and operating correctly.

**TESTING ABR
CATALOG
LOCATE EXIT**

Run batch jobs and execute TSO commands and ISPF functions which reference data sets which have been ARCHIVED for automatic recall. Make sure not to specify volume serial information, which causes the catalog to be bypassed. Verify that the referenced data sets are automatically restored.

**TESTING ABR
DATA SET NOT
FOUND EXIT**

Run batch jobs that OPEN data sets which have been archived, making sure to specify unit and volume serial information (UNIT= and VOL=SER=) in the DD statement, so that the ABR CATALOG LOCATE exit will not recall the data sets, and the ABR Data Set Not Found (DSNF) exit will be invoked. Verify that the OPENED data sets are automatically restored.

For non-VSAM data sets, IEBGENER can be used to read the data. For VSAM, IDCAMS can be used to PRINT or REPRO a cluster.

Warning: contrary to what many believe, IEFBR14 does not OPEN the data sets in its JCL, so IEFBR14 is not an appropriate utility to test the DSNF exit (with UNIT= and VOL= specified in the JCL). However, IEFBR14 can be used to test the LOCATE exit (omitting UNIT= and VOL=).

**TESTING ABR
DADSM PRE-
PROCESSING
EXIT**

SCRATCH data sets which current ABR backups (you may need to run ABR full-volume or incremental backups first). Verify that the SCRATCHed data sets are recorded in the ABR SCRATCH catalog: either by using the ISPF dialog for the ABR SCRATCH Catalog Report (ISPF option A.1.3), or by using program FDRABRP with the PRINT SCRATCH command.

RENAME data sets which have been backed up by ABR. Verify that the old data set names are recorded in the SCRATCH catalog, as above. Also verify that the renamed data sets do not appear to have a current ABR backup: either by using the ISPF dialog for the ABR BACKUP Report (ISPF option A.1.2), or by using program FDRABRP with the PRINT BACKUP command.

**TESTING
FDRREORG
IEBCOPY
INTERCEPT**

Run an IEBCOPY compress of a PDS and verify that the SYSPRINT output contains FDR messages, instead of IEBCOPY messages. Include a STEPLIB DD pointing to the FDR load library unless it is in the LINKLIST. Next, include an //NFDRCOPY DD DUMMY in the JCL stream and rerun the job to verify that IEBCOPY is successfully invoked.

**TESTING A
NEW VERSION**

If the Operating System exits from a prior version of FDR are dynamically installed in your system in production, it is possible to test a new version of the exits, or a change in options, for a specified set of jobs while continuing to use the production version of the exits for other jobs.

The procedure to do this is as follows:

TEST LIBRARY

Create an APF-authorized test FDR program library, if you don't already have one. This library may be any of the following:

- a complete library for a new release of FDR/ABR.
- a complete copy of the library for the current release of FDR/ABR.
- a partial library for either a new release or the current release. This library must contain at least the following modules: FDRSTART, FDROPT, FDRVECTB, and whichever exit modules (FDR026DU, FDREXDSN, FDRPRE00, FDRBCOPY, FDRDAS02 and/or FDR00024) will be dynamically installed for this test.

90.62 CONTINUED

**CATALOGED
PROCEDURE**

If you will be testing the CATALOG LOCATE exit, you will probably need to copy the 'SYNRECAL' procedure from the FDR Installation Control Library (ICL) for the release being tested to a cataloged procedure library that is available to JES for START commands; since you probably already have a SYNRECAL procedure in place, you should rename this member while copying it. Edit this member and specify the dsname of the test library on the STEPLIB DD statement so that recalls done by the procedure will use the proper version of ABR. You will need to use ISPF Panel A.I.4.11.2 to change the LXSYNPROC option to the new procedure name, in the copy of FDROPT in the test library.

If you are testing only new or modified exits, using the production version of ABR, and no option changes are required, the above is probably not necessary; the production version of the SYNRECAL proc can be used and the LXSYNPROC name change is not required.

**SETTING TEST
OPTIONS**

Set the desired options in the test library. Be sure to specify the name of the test library on panel A.I.4. If you are testing an option change, then of course some options in the test library will be different from the production version. If you are testing a new version of the exits, then all of the options may be the same as for the production version, and this step may not be necessary.

If you are testing a new version of a particular exit, or a change in an option that affects only one exit, then it is appropriate to set the options in the test library to dynamically install only that exit. For example, you could set ABRLOC to YES and set the other dynamic installation options to NO in the test library, even though all options are set to YES in the production Option Table. On the other hand, if you are testing an exit that you have never used before, you would set the option to install that exit to YES in the test library even though it is NO in the production library in LINKLIST.

**INSTALL FOR
TEST**

Dynamically install the test version of the exits by executing a batch job using the following JCL:

```
//FDRSTR TT JOB 'DYNAMICALLY INSTALL TEST VERSION OF EXITS',FDR
//FDRSTR TT EXEC PGM=FDRSTART,
// PARM='TEST,JOBNAME=jobname' See Note 1
//STEPLIB DD DSN=fdrabr.test.library,DISP=SHR See Note 2
//SYSLIB DD DSN=fdrabr.test.library,DISP=SHR See Note 2
//ABNLDUMP DD DUMMY FOR ABEND-AID
//SYSUDUMP DD SYSOUT=A
```

This JCL is supplied in member FDRSTR TT in the Installation Control Library (ICL). You can also generate this job with the REFRESH command on Panel A.I.4 (the TEST parm can be specified on the FDRSTART panel).

Note 1: On the EXEC statement, JOBNAME= specifies the job(s) that are to be processed by the test version of the exits. JOBNAME= may be abbreviated as JOB=. The value specified may be:

- 1) a single jobname, e.g., JOBNAME=XYZ
- 2) a jobname prefix followed by an asterisk, e.g., JOBNAME=XYZ*
- 3) up to 5 jobnames or prefixes enclosed in parenthesis, e.g., JOB=(ABC,XYZ*,SYS*). This multiple jobname format cannot be used if the production version of the exits is less than V5.2 level 51.

TSO userids and started task names will also be checked for a match and will be eligible for processing by the test version of the exits. Any job that does not match will be processed by the production version of the exits. If JOBNAME=* or JOB=* is specified, then all jobs will be processed by the test version of the exits; in effect, the production version of the exits will be temporarily replaced.

Note 2: The SYSLIB and STEPLIB DD statements must specify the dsname of the test FDR program library.

Note 3: The FDRSTATS JCL shown in [Section 90.61](#) will display the status of both the production and test exits.

90.62 CONTINUED

**ABRLIB/
STEPLIB FOR
TEST**

If you are testing a new version of the CATALOG LOCATE exit and/or the Data Set Not Found (DSNF) exit, any auto-recalls must use the programs in the test FDR program library. Since auto-recalls may occur in the address space of the caller:

-- the JCL for batch jobs and started tasks using the test exits (specified by JOBNAME= described above) must be modified to include this DD statement in any step that may do an auto-recall:

```
//ABRLIB DD DSN=SHR,DSN=test.library.name
```

-- the logon procedure for TSO users which use the test exits may need to be updated to include the same ABRLIB DD statement, or a TSO ALLOC command must be issued for the same DDname and library before any recalls are done.

The test program library should contain a copy of FDR from the same release as the exits being tested.

Note: although a copy of FDROPT is loaded from the test program library by FDRSTART when the test exits are installed, it is available only to the exits themselves. Any programs running under the test exits will get FDROPT from the test library directly via the ABRLIB or STEPLIB DD statements.

**UPDATING
TEST OPTIONS**

Once test exits is installed, you must specify PARM=TEST or PARM=PROD for every execution of FDRSTART until the test exits are removed by the next IPL, to clearly indicate which exits you are modifying.

To change the options used by the test exits:

Change the desired options in the test library, using the ISPF panels and execute a batch job, using the following JCL:

```
//FDRCHOPT JOB 'CHANGE OPTIONS FOR TEST ABR EXITS',FDR
//FDRSTART EXEC PGM=FDRSTART,PARM=TEST
//STEPLIB DD DSN=fdrabr.test.library,DISP=SHR
//SYSLIB DD DSN=fdrabr.test.library,DISP=SHR
//ABNLDUMP DD DUMMY FOR ABEND-AID
//SYSUDUMP DD SYSOUT=A
```

FDRSTART will replace the copy of FDROPT used by the test version of the exits with the copy that it reads from the test library (pointed to by the SYSLIB DD statement). If that option table enables dynamic installation of any exits that were not previously installed for test, they will be installed.

Since the JOBNAME parameter was not specified, the test will continue for the same job name(s) as were specified on the previous FDRSTART. If necessary, repeat this on every system where the test exits are active.

To change the jobname(s) used by the test exits:

Execute a batch job, using the following JCL:

```
//FDRCHOPT JOB 'CHANGE JOBNAMES FOR TEST ABR EXITS',FDR
//FDRSTART EXEC PGM=FDRSTART,PARM='TEST,JOBNAME=jobname'
//STEPLIB DD DSN=fdrabr.test.library,DISP=SHR
//SYSLIB DD DSN=fdrabr.test.library,DISP=SHR
//ABNLDUMP DD DUMMY FOR ABEND-AID
//SYSUDUMP DD SYSOUT=A
```

The test version of the exits will now process the job(s) specified by the new FDRSTART PARM value. The job(s) that were previously being processed by the test version of the exits will now be processed by the production version of the exits. FDRSTART will also replace the copy of FDROPT used by the test version of the exits as described above; any changes in the options will become effective.

You can effectively deactivate the test exits by specifying a jobname which is meaningless (will never match any real jobname, TSO userid or started task name).

90.62 CONTINUED

To update the options in use by the production exits while a test exit is active:

Execute FDRSTART as shown in [Section 90.61](#) Step 7, except that the EXEC statement must specify:

```
//FDRSTART      EXEC      PGM=FDRSTART , PARM=PROD
```

**TESTING ONLY
OPTION
CHANGES**

You can use the test exit procedure to test changes in the FDR Global Option Table with the production version of the exits, to test the effect of option changes which affect the operation of the exits. To do so easily:

- 1) update the required options, using the ISPF panels. You can do this in the production FDR program library since it will not affect the copy of FDROPT in use by the production exits. No test program library is required.
- 2) active the "test" exits, which will really be the production exits with the updated copy of FDROPT. You can omit the STEPLIB DD statement from the FDRSTART jobs, and you will not need an updated copy of the SYNRECAL procedure, nor will you need to use ABRLIB or STEPLIB. Only the jobs and users specified by the JOBNAME= parameter will use the new options.
- 3) when testing is complete, you can activate the new options in production with the FDRSTART PARM=PROD as shown above. Or, if you decide not to put the new options in production, return the options in the FDR program library to their original values (so that the next IPL will not pick up the modified FDROPT).

90.70 APPLYING FDR FIXES (FDRSPZAP)

Fixes to the FDR products will usually be provided by Innovation in the form of zaps. These zaps can be applied using the IBM SUPERZAP program (PGM=AMASPZAP or IMASPZAP) but Innovation also provides a program for applying zaps which has several advantages.

FDRSPZAP The FDRSPZAP program uses JCL almost identical to the IBM SUPERZAP program and accepts all SUPERZAP control statements. But it has 2 advantages over SUPERZAP:

- The CHECKSUM statement, which Innovation provides as part of all zaps, is used as a check that zaps have been entered correctly. It compares the value provides to a checksum of all hexadecimal digits on the preceding VER and REP statements; almost all typing errors will result in a checksum mismatch and an error message. However, the IBM SUPERZAP program compares the checksum AFTER it has already applied the changes in the REP statements. So it tells you that you have just applied a bad zap and leaves it to you to back it off. FDRSPZAP will read all the control statements and verify the CHECKSUM before invoking SUPERZAP to perform the actual zap.
- FDRSPZAP can apply zaps to the in-storage copies of the FDR system exits. If test exits are installed, fixes can be applied to either the test or production exits.

FDRSPZAP JCL The JCL for execution of FDRSPZAP will look like this (similar to SUPERZAP):

```
//FDRSPZAP    EXEC      PGM=FDRSPZAP , PARM=parm
//STEPLIB     DD        DISP=SHR,DSN=fdr.program.library
//SYSLIB      DD        DISP=SHR,DSN=fdr.program.library
//SYSPRINT    DD        SYSOUT=*
//SYSIN       DD        *
...SUPERZAP Control Statements supplied by Innovation...
/*
```

The STEPLIB DD can be omitted if the FDR program library is in the system linklist. However, zaps to a level of FDR still being tested should be applied with the FDRSPZAP from the test library.

If the zap being applied affects one of the dynamic installed FDR exit modules, the in-storage module is automatically updated as well.

FDRSPZAP PARM FDRSPZAP accepts a PARM= value ("parm" in the example JCL above). Possible values are:
null or PARM= omitted - FDRSPZAP applies the zap only to the library indicated by SYSLIB.

LPA - for zaps to an FDR dynamic exit, the zap is applied to the in-memory copy only.

BOTH - zaps are applied to the library indicated by SYSLIB; for zaps to an FDR dynamic exit, they are also applied to the in-memory copy.

PROD - if both test and production exits are installed, PROD is required to apply zaps to the in-memory copy of production exits.

TEST - if both test and production exits are installed, TEST is required to apply zaps to the in-memory copy of test exits.

IGNIDRFULL - all Innovation zaps include an IDRDATA statement, which stores a string with each module indicating which zaps are applied. If the space reserved for IDR data in a load module becomes full, IGNIDRFULL allows the zap to be applied anyway (although the IDRDATA is not saved).

Multiple values can be combined, in parenthesis. For example, to apply zaps to test exits you must specify PARM=(LPA,TEST) or PARM=(BOTH,TEST)

THIS PAGE WAS INTENTIONALLY LEFT BLANK

91.01 FDR ALTERNATE INSTALLATION PROCEDURE

Innovation strongly recommends the use of the interactive installation procedure documented in [Section 90](#). It is easier to use and less prone to error than any manual procedure. The interactive procedure must be used if your distribution tape includes ABR. But if you cannot use ISPF or prefer not to do so for a non-ABR installation, you can follow the alternate procedure in this section to install FDR.

You must still follow the checklist of install steps shown in [Section 90.01](#), but the procedures for executing those steps will change:

- You can use the manual procedure described in [Section 91.02](#) to load the FDR distribution libraries to disk.
- The FDRZAPOP utility ([Section 91.03](#)) will be used for setting all FDR Global Options (except FDRREORG options) and for defining the COMPAKTOR Unmovable Table. You will need to read the sections in [Section 90](#) referenced by the checklist in order to understand the options, tables, and lists, but you will use FDRZAPOP to actually do the work.
- The FDRREOZO utility ([Section 31](#)) will be used for setting all FDRREORG options and defining the FDRREORG NOREORG List.
- If you will be using the FDR ISPF dialogs at all, FDR dialog options must be set by editing the ABRGLOBL member in the FDR CLIST library.
- SAR can be installed on disk and/or tape using the FDRSARLR program described in [Section 15.20](#).

91.02 LOADING THE FDR LIBRARIES

If you do have access to TSO, Innovation recommends that you use the Interactive FDRLOAD program described in [Section 90.03](#) to load the FDR libraries from the distribution tape to disk.

But, if you are unable to use FDRLOAD, the following procedure can be used to load the libraries.

First, copy the install jobstream from the ICL (Installation Control Library) on tape to disk, using JCL similar to:

```

JCL FOR      //INSTALL      JOB      ...
LOADING     //COPY        EXEC    PGM=IEBCOPY
INSTALL     //SYSPRINT    DD      SYSOUT=*
JOBSTREAM  //TAPEIN      DD      DSN=ICL,DISP=OLD,LABEL=4,UNIT=tape,
//              VOL=SER=FDR53T    <-- Change to FDR53P if production tape
//LIBRARY      DD      DSN=your.library.here,
//              DISP=SHR
//SYSIN        DD      *
              COPY INDD=((TAPEIN,R)),OUTDD=LIBRARY
              SELECT MEMBER=INSTALL
//*
```

You will need to specify an appropriate unit name for the tape; if your tape is a production tape, be sure and change the tape volume serial to FDR53P. The LIBRARY DD statement should point to a control card library (RECFM=FB and LRECL=80); this can be an existing library or you can allocate a new one. The jobstream above will load **only** the single member **INSTALL**.

If you want to **only** install the FDR Load Library you can use the member **FDRLINK**.

**INSTALLA-
TION JCL**

The JCL member loaded by the preceding job provides you with procedures for the installation of the FDR:

- Product load library
- Installation Control library (ICL)
- JCL (example) library
- ISPF dialog libraries

The JCL stream consists of three in-stream procedures and the necessary steps to execute those procedures to install the FDR libraries. One procedure allocates the target libraries and contain the recommended size of each library; the other two execute IEBCOPY or IEBUPDTE to load the library from tape to disk. The jobstream contains comments which will guide you through modifying the jobstream to substitute data set names and target disk volumes of your choice. Once modifications are complete, submit the jobstream to allocate and load the libraries.

Note: The FDR load library should be an authorized program library as shown in [Section 90.40](#).

91.03 EXECUTING FDRZAPOP

Sections 90.10 through 90.22 of the interactive FDR install procedures describe use of the FDR dialogs to change the permanent defaults for many FDR options. Section 90.30 describes use of the dialogs for creating the COMPAKTOR Unmovable Table. If you cannot use the dialogs for these functions, you can perform all of the same functions with FDRZAPOP, the FDR Global Option utility, executed as a batch job.

Options for FDRREORG and the FDRREORG NOREORG List (Sections 90.23 and 90.31) are handled by FDRREOZO, the FDRREORG Global Option Utility, described in Section 31.

This section describes use of FDRZAPOP, but it contains only brief descriptions of the commands available. Extensive HELP descriptions can be printed or displayed to help you use FDRZAPOP; member FDRZAPOP of the FDR ICL (Installation Control Library) contains a complete printout of the FDRZAPOP HELP which can be printed or browsed.

You will still need to go through Sections 90.10 to 90.30 to understand the options and tables/lists you need to set or create, but you will use FDRZAPOP to perform the changes. To simplify the process, the PRINT command of FDRZAPOP will print the current values of the FDR options in the same organization as the FDR dialog panels described in those sections; even the panel numbers are referenced.

FDRZAPOP JCL To execute FDRZAPOP as a BATCH job use the following JCL:

```
//ZAPOP EXEC PGM=FDRZAPOP
//STEPLIB DD DISP=SHR,DSN=fdr.program.library (if required)
//SYSPRINT DD SYSOUT=*
//SYSLIB DD DISP=SHR,DSN=fdr.program.library
//SYSIN DD *
... FDRZAPOP Control Statements ...
```

**COPY
OPTIONS
FROM
PREVIOUS
VERSION**

When you upgrade to a new release of FDR, you will usually want to copy all FDR options from the release currently in use. You do so with the COPY command of FDRZAPOP. This sample jobstream will copy all FDR option settings as well as the ABR protect lists, the ABR Restore Allocation list, and the COMPAKTOR unmovable table.

```
//COPY EXEC PGM=FDRZAPOP
//STEPLIB DD DISP=SHR,DSN=new.fdr.program.library
//SYSPRINT DD SYSOUT=*
//SYSLIB DD DISP=SHR,DSN=new.fdr.program.library
//OLDFDR DD DISP=SHR,DSN=old.fdr.program.library
//SYSIN DD *
COPY ALL,FROMDD=OLDFDR
```

**FDRZAPOP
COMMANDS**

The FDRZAPOP supports the following commands (see the HELP for details):

- CANCEL** -- Terminates FDRZAPOP without saving any option changes. Used under TSO.
- COPY** -- copies options and tables/lists from a previous version of FDR.
- CPKUNMOV** -- Add entries to the COMPAKTOR Unmovable Table.
- END** -- Terminates FDRZAPOP and saves any changed options. Used under TSO. In batch an END command is executed automatically when the end of the control statements is reached (EOF on SYSIN).
- HELP** -- The HELP command will print a menu of the FDRZAPOP options and related documentation.
The HELP command format is:
HELP or HELP SUB(command) or HELP ALL
- PRINT** -- If the PRINT command is specified without any operands, the current values in the FDR Global Options Table will be printed. If one or more table/list names are specified as operands on the PRINT command, the specified tables will be printed.
- RESET** -- If the RESET command is specified without any operands, the FDR Global Options Table will be reset to the original values on the installation tape. If one

91.03 CONTINUED

or more table/list names are specified as operands on the RESET command, the specified tables will be reset to their original contents on the installation tape. An operand of ALL will reset all options and all tables and lists.

ZAP -- Modify options in the FDR Global Options Table. This command enables the installation to enable or disable specified functions and set control statements defaults and other processing options for the FDR system.

The ZAP command format is:

**ZAP ENABLE=(option1,...,optionn)
 ,DISABLE=(option1,...,optionn)
 ,option1=value,...,optionn=value**

100.01 INTRODUCTION TO MESSAGES

This section describes the various printer, console and TSO messages which may be output by Innovation's FDR product line and the various ABEND codes with which FDR programs may terminate.

FDR MESSAGE FORMAT

Except for FDRREORG and FDRCPK, messages from all components of FDR have this format:

FDRnnn Message-text

where "nnn" is a 3-digit message number. When FDRnnn is followed by ** (2 asterisks) the message usually indicates some sort of error, which may result in an ABEND or a non-zero return code. Messages without the asterisks are usually informational. These messages are described in [Section 100.03](#), [100.04](#) (for SAR), and [100.05](#).

Messages directed specifically to the MVS console or TSO user will have this format:

FDRWnn Message-text

Some of these messages require replies by the MVS operator or TSO user. They are documented in [Section 100.02](#). However, other FDRnnn messages may also be sent to the MVS console or TSO user.

CPK MESSAGE FORMAT

The general format of COMPAKTOR messages is as follows:

CPKnnns Message-text

where "nnn" is a 3-digit message number and "s" is a severity code:

- I Informative message.
- A Action required message.
- W Warning message. May result in abnormal termination.
- E Error message. Always results in abnormal termination.

Some of these message may also be issued to the MVS console. COMPAKTOR also produces FDRnnn messages for some conditions. COMPAKTOR messages are documented in [Section 100.06](#)

**FDRREORG MESSAGE
FORMAT**

The general format of FDRREORG messages is:

FDRRnn Message-text

FDRSnn Message-text

where "nn" is a 2-digit message number. If the message number is followed by * (single asterisk) this is a warning message, while if it is followed by ** (double asterisk) it is an error message. FDRREORG messages are documented in [Section 100.07](#).

RETURN CODES

FDR sets a return code at the end of the step, unless it ABENDs:

0 - normal completion

4 - a trial version of FDR has reached the end of its trial period and will no longer function. If you have licensed a production version of FDR you should install it. Contact Innovation if you wish to continue testing with the trial version. A return code of zero indicates that FDR has performed all functions successfully. A return code of 4 indicates that the testing period for the product has expired.

Any other return code - errors of some kind have occurred during this execution of FDR. Check the listing for the error messages. This usually indicates that errors occurred which were not severe enough to terminate the FDR operation, yet FDR wanted to call attention to the error messages at the end of processing. Severe errors usually result in an immediate ABEND.

You can control whether FDR programs will set a return code or issue a U0888 ABEND for non-terminating errors; see the options FDRCC, CPKCC, and ABRCC in [Section 90](#). By default, FDR and COMPAKTOR will issue a U0888 ABEND and ABR will set a return code 12.

ABEND CODES

FDR user ABEND codes range from U0100 to U0999 and are documented in [Section 100.08](#). Most user ABENDs are preceded by an FDR error message.

Note that user ABEND U0888 is a special case. It is issued by FDR programs to indicate that errors occurred which were not severe enough to terminate the FDR operation, yet FDR wanted to call attention to the error messages at the end of processing. Severe errors usually result in an immediate ABEND.

You can control whether FDR programs will set a return code or issue a U0888 ABEND for non-terminating errors; see the options FDRCC, CPKCC, and ABRCC in [Section 90](#). By default, FDR and COMPAKTOR will issue a U0888 ABEND and ABR will set a return code 12.

WAIT STATE CODES

The Stand Alone Restore (SAR) program may put the CPU in a WAIT STATE if errors occur and it is not able to display a message to the operator, e.g., errors that occur before the console has been initialized. SAR will load a meaningful code in the instruction address field of the PSW (last 4 digits); these codes are documented in [Section 100.09](#). The procedures for displaying the PSW vary by CPU type; consult the hardware documentation for your CPU for details.

100.01 CONTINUED . . .

MINI-DUMP

In many cases, an error message will be followed by a set of diagnostic displays, called an FDR mini-dump. Each message description indicates if a mini-dump is printed. The mini-dump includes:

- the general registers at the time of the mini-dump. These may or may not be significant, depending on the error.
- blocks of storage identified by the FDR program requesting the mini-dump. These usually have a title above them for identification. The storage displayed varies depending on the error message.
- for I/O errors, it will display the DCB, UCB, and IOB. The IOB contains:

remaining length in last CCW (CCW length field minus this gives bytes transferred)

bytes 2-3	sense data
byte 4	I/O termination post code (7F = normal termination)
bytes 8-15	hardware-generated CSW (channel status word, see the IBM Principles of Operation manual), consisting of:
bytes 8-11	address of last CCW executed +8
bytes 12-13	channel status flags
bytes 14-15	remaining length in last CCW (CCW length field minus this gives bytes transferred)

It will also format the CCW chain executed. CCW(-0) identifies the last CCW executed.

100.02 FDR CONSOLE MESSAGES AND TSO TERMINAL MESSAGES

The following write-to-operator (WTO), write-to-operator-with-reply (WTOR) and TSO messages are issued by FDR:

- FDRW01 CONFIRM REQUEST TO RESTORE UNIT=uuuu,VOL=vvvvvv,JOB=jjjjjjj**
Reason: The job 'jjjjjjj' wishes to do a full-volume restore or copy to the disk device 'uuuu' whose current volume serial is 'vvvvvv'. This confirmation request can be suppressed by specifying the CONFMESS=NO operand on the RESTORE or COPY statement.
Action: Reply YES--Operator approves request to restore.
 Reply NO-- Operator does not approve the restore request; the restore is terminated
- FDRW02 CONFIRM REQUEST TO ABSOLUTE TRACK RESTORE UNIT=uuuu, VOL=vvvvvv, JOB=jjjjjjj**
Reason: The job 'jjjjjjj' wishes to do an absolute track restore to volume 'vvvvvv' on the disk device 'uuuu'.
Action: Reply YES--Operator approves request to restore.
 Reply NO-- Operator does not approve the restore request; the restore is terminated
- CONFIRM REQUEST TO CONTINUE BACKUP - PREVIOUS INSTANT BACKUP FOR vvvvvv WILL BE LOST - REPLY YES OR NO**
Reason: This form of FDRW02 is issued by FDR InstantBackup. You have done a SNAP/SPLIT operation for volume "vvvvvv" but you have not yet executed the ABR DUMP with SNAP=USE or BCV=USE to move the backup to tape, and you are trying to execute another SNAP/SPLIT. If you specify the operand CONFMESS=NO on the SNAP/SPLIT statement, then FDR will overlay backup image on disk unconditionally, and issue message FDR230 REASON=P.
Action: Reply YES--Operator approves request to backup. The backup already on the SNAP target volume or BCV will be lost.
 Reply NO-- the backup will be terminated.
- INSTANT BACKUP FOR vvvvvv ON SNAPUNIT=uuuu IS IN-USE. REPLY "RETRY" OR "CANCEL"**
Reason: This form of FDRW02 is issued by FDR InstantBackup. You have done a SNAP/SPLIT operation for volume "vvvvvv" but the ABR DUMP with SNAP=USE or BCV=USE to move the backup to tape or a restore job is still executing, and you are trying to execute another SNAP/SPLIT.
Action: Reply RETRY—when you believe the backup will no longer be in use. If it is still busy, the message is reissued.
 Reply CANCEL-- the backup will be terminated.
- FDRW03 dsname DATE PROTECTED--REPLY YES OR NO TO RESTORE VOL=vvvvvv**
Reason: A data set restore has been requested for data set "dsname" on volume "vvvvvv". The output data set is pre-allocated and has an expiration date that has not been reached. If you wish to suppress this message and automatically do the restore, specify EXPD=NONE on the RESTORE statement.
Action: Reply YES--Operator approves request to restore.
 Reply NO-- Operator does not approve the restore request; the restore is terminated
- FDRW04 FULL VOLUME RESTORE COMPLETED--UNIT=uuuu SHOULD BE REMOUNTED ON OTHER SYSTEMS**
Reason: A full volume restore was completed on the disk device "uuuu" which is marked as "Shared DASD" in your I/O configuration. FDR has determined that the restore changed the volume serial and/or the VTOC location on the output volume. FDR has made the necessary changes on the system where the restore was run, but if other CPUs access the disk, they will not know of this change and may not process this volume properly.
Action: Normally all other systems should make this device offline prior to the restore, especially if the volume serial or VTOC will change. This message indicates that you may now remount the volume on the other systems; vary it OFFLINE if not already offline, then vary it ONLINE.
- FDRW05 DISK UNIT=uuuu VOL=vvvvvv HAS BEEN RELABELED TO VOL=nnnnnn**
Reason: Disk device "uuuu" has been restored. Its original volume serial was "vvvvvv" but the restore has changed the volume serial to "nnnnnn".
- FDRW06 DISK UNIT=uuuu VOL=vvvvvv NOW OFFLINE--DUPLICATE VOLSER**
Reason: Disk device "uuuu" has been restored. The restore has changed the volume serial to "vvvvvv", which matches the volser of another online volume (probably the original volume).
Action: FDR will set disk device "uuuu" offline.

100.02 CONTINUED . . .

- FDRW10 FDRDRP COULD USE SOME ADDITIONAL type TAPE DRIVES. PLEASE VARY SOME ONLINE IF AVAILABLE.**
Reason: A FDRDRP restore has not reached its limit of active tape drives (MAXTAPES=) but no additional drives of the needed type are online.
Action: Vary one or more tape drives online. After a few moments, FDRDRP will begin to use them.
- FDRW22 DSN *dsname* PASSWORD PROTECTED**
Reason: Data set "*dsname*" is password protected and is being backed up or restored. This refers to old OS password protection for non-VSAM data sets and VSAM password protection for ICF VSAM. It does not refer to security systems like RACF.
Action: To continue processing the operator must reply to the appropriate MVS password message.
- FDRW23 DSN *dsname* DATE PROTECTED REPLY YES OR NO FOR ARCHIVE**
Reason: Non-VSAM data set "*dsname*" was selected for Archive or Superscratch, but the data set has an expiration date which has not been reached. If you wish to suppress this message and automatically select the data set, specify EXPD=NONE on the DUMP statement.
Action: Reply YES--Archive and/or scratch the data set.
Reply NO--DO NOT select the data set.
- FDRW24 FDRABR REQUEST FOR TAPES--vvvvvv... REPLY YES NO OR BYPASS**
Reason: The OPERATOR operand was specified on a RESTORE statement. A restore of a data set was requested from the backup volumes "vvvvvv...".
Action: Reply YES--You have the volumes available.
Reply NO--You do not have the volumes available, or do not wish to mount them at this time. The restore is cancelled.
Reply BYPASS--same as NO.
- FDRW25 ABR REQUEST FOR TAPES(*nn*) vvvvvv... REPLY YES NO OR BYPASS(*NN*)**
Reason: The OPERATOR operand was specified on a RESTORE statement for a full-volume ABR RESTORE. Message FDRW25 will be repeated for each ABR Volume Backup that will be used during the restore; each line shows the ABR Cycle number (*nn*) followed by the volume serials of the backups ("vvvvvv...") required for that cycle.
Action: Reply YES--All of the volumes are available.
Reply NO--All or some volumes are not available. The restore request is cancelled.
Reply BYPASS(*nn*)--The tapes shown for cycle "*nn*" will not be used during the restore. "*nn*" may be any cycle that was listed, except cycle zero (the full volume backup). After a reply of BYPASS(*nn*), the message will be reissued so that additional backups can be bypassed. When all desired backups to be bypassed have been specified, reply YES. Data sets whose latest backup was on the cycles that you bypassed will probably not be correctly restored.
- FDRW26 INVALID RESPONSE REPLY YES NO OR BYPASS(*NN*)**
Reason: The operator reply to Message FDRW25 was invalid. Either, the operator reply was an invalid keyword, the cycle number was invalid or did not exist, or the cycle number was zero.
Action: Give a proper response to the message.
- FDRW27 REPLY WAIT NOWAIT OR RETRY FOR ENQ WAIT ON DSN=*dsname***
Reason: An FDR job specified DSNENQ=HAVE. Data set "*dsname*" is currently ENQed to another job.
Action: Reply WAIT -- FDR will wait for the data set to become available. The job could time out if the ENQ is not released in a short time.
Reply NOWAIT -- FDR will issue a FDR158 warning message for the data set.
For DUMP the data set will be dumped unless ENQERR=BYPASS is specified.
For Archive or Superscratch the data set will be bypassed.
For RESTORE the data set will be bypassed.
For COMPAKTOR the data set will be made unmovable.
Reply RETRY -- FDR will retry the ENQ. If it is still not available, message FDRW27 is issued again.
- FDRW40 MODULE *modname* INSTALLED AT *address* -- VER *nn***
Reason: The message is issued in response to a status request when the FDR operating system exit module "*modname*" is installed in the system. "*address*" is the virtual storage location of that module. The version number of the module is also given.

100.02 CONTINUED . . .

- FDRW41 THE FDR SYSTEM MODULES ARE *status -- VER nn type***
Reason: This message is issued by FDRSTART or FDRSTATS when the FDR operating system exit modules are installed in the system. The "status" of the modules may be:
ACTIVE on the first execution of FDRSTART or on executions of FDRSTATS;
RE-ACTIVATED on executions of FDRSTART where an FDR operating system exit module was re-activated.
REFRESHED on subsequent executions of FDRSTART where the copy of the FDR Global Options Table module in global storage was refreshed.
This message includes the version and release level of the modules installed. "type" may indicate
PROD for production exits
TEST for test exits (the JOBNAME= or JOBGROUP= that will use the test exists is also displayed).
- FDRW42 THE FDR SYSTEM MODULES ARE *status type***
Reason: This message is issued by FDRPARE or FDRSTATS when the FDR operating system exit modules are disabled. The "status" of the modules may be:
NOT ACTIVE when the ABR operating system exit modules are not installed in the system.
INACTIVE when the ABR operating system exit modules are installed in the system but are disabled, i.e., not processing any requests.
"type" may indicate
PROD for production exits
TEST for test exits (the JOBNAME= or JOBGROUP= that will use the test exists is also displayed).
- FDRW44 *module* ESTAE RECOVERY ENTERED FOR ABEND *Ssss Uuuuu* AT OFFSET *xxxx***
Reason: This message is issued by "module", which can be FDRSTART, FDRSTATS or FDRPARE, when an unexpected system (sss) or user (uuuu) ABEND occurs at displacement "xxxx" in the module. If the offset is not displayed, the abend did not occur within the FDRSTART module.
Action: **If the problem persists call Innovation for assistance.**

100.02 CONTINUED . . .

- FDRW46** **module FAILURE -- REASON=reason**
Reason: The execution of "module", which can be FDRSTART, FDRSTATS or FDRPARE, failed for one of the following reasons:
- 01 -- OPERATING SYSTEM NOT MVS SP 1.2 OR HIGHER.**
 - 02 -- INVALID INPUT PARAMETERS.** Check the PARM= specified.
 - 03 -- UNABLE TO OBTAIN AUTHORIZATION.** Program library must be APF authorized.
 - 04 -- FDRVECTB VECTOR TABLE IS INVALID.** Call Innovation for assistance.
 - 05 -- SYSLIB DD MISSING OR OPEN ERROR.** Check the SYSLIB DD statement.
 - 06 -- modname NOT FOUND IN SYSLIB.** The ABR exit "modname" was not found in the library pointed to by the SYSLIB DD statement.
 - 07 -- modname MODULE LOAD ERROR.** Call Innovation for assistance.
 - 08 -- ERROR MODIFYING THE SVC TABLE.** Call Innovation for assistance.
 - 09 -- modname NOT FOUND.** The module "modname" was not found in the link pack area. Call Innovation for assistance.
 - 10 -- modname NOT FOUND IN THE LINKLIST.** The ABR exit module "modname" was not found in an APF authorized LINKLIST library.
 - 11 -- modname UNABLE TO DE-INSTALL.** Call Innovation for assistance.
 - 12 -- modname CDE or LPDE ABOVE 16M ERROR.** Call Innovation for assistance.
 - 13 -- modname MODULE ABOVE 16M ERROR.** Call Innovation for assistance.
 - 14 -- GETMAIN ERROR.** Increase the region in the step JCL.
 - 15 -- modname HAS SMP INSTALLED IDP MODULE.** FDRSTART detected the presence of an ABR operating system exit installed via SMP in the named module. The ABR operating system exits that are SMP installed must be removed if they are to be dynamically installed.
 - 16 -- UNABLE TO OBTAIN CMS LOCK.** Call Innovation for assistance.
 - 17 -- FREEMAIN ERROR.** Call Innovation for assistance.
 - 18 -- UNABLE TO OBTAIN LOCAL LOCK.** Call Innovation for assistance.
 - 19 -- modname CDE NOT FOUND.** Call Innovation for assistance.
 - 21 -- ENQUEUE/DEQUEUE ERROR ON SYSZSVC.** Call Innovation for assistance.
 - 22 -- INVALID VECTOR TABLE STATUS.** Call Innovation for assistance.
 - 23 -- INVALID VECTOR TABLE CDE ADDRESS.** Call Innovation for assistance.
 - 24 -- RC=xxxx FROM SVCUPDTE.** Call Innovation for assistance.
 - 25 -- NO MODULES TO INSTALL -- FDROPT.** The dynamic installation options that indicate which of the ABR operating system exits are to be installed specified that none of the exits are to be installed. These options (ABRLOC, ABRDSNF, ABRPRE, IEBCOPY) are set in the FDR Global Options Table, module FDROPT, using the ABR Install dialog (ISPF panel A.1.4.11.1).
 - 26 -- FDR SYSTEM MODULES INSTALL CANCELLED.** This message is displayed when the console operator replies NO to the FDRW47 message.
 - 29 -- FDROPT MODULE LEVEL IS NOT HIGH ENOUGH.** The dynamic installation options are not present in the level of FDROPT that was loaded.
 - 30 -- PARM (PROD or TEST) NOT SPECIFIED.**
 - 32 -- modname ERROR.** Call Innovation for assistance.
 - 33 -- modname ERROR.** Call Innovation for assistance.
 - 34 -- SYSEVENT DONTSWAP ERROR.** Call Innovation for assistance.
 - 35 -- PROGRAM LEVEL LOWER THAN VECTOR TABLE.** Call Innovation for assistance.
 - 36 -- VECTOR TABLE TO BE REPLACED NOT FOUND.** Call Innovation for assistance.
 - 37 -- TBL LEVEL NOT HIGH ENOUGH FOR REPLACE.** Call Innovation for assistance.
 - 38 -- modname - CSVDYLP RC=xxxx-xxxxxxxx.** Call Innovation for assistance.
 - 39 -- modname NOT FOUND.** Call Innovation for assistance.
 - 40 -- modname ENQUEUE ERROR ON SYSZCSV/CSVDYLP.** Call Innovation for assistance.
- Action:** If error cannot be determined, contact Innovation for assistance.
- FDRW47** **CONFIRM REQUEST TO INSTALL THE FDR SYSTEM MODULES**
Reason: This WTOR prompts the console operator to confirm installation of the FDR operating system exits. This message is issued unless the FDR Global Option CONFINST is set to NO.
Action: Reply 'YES' to proceed or 'NO' to bypass the dynamic installation of the FDR operating system modules.
Note: The FDR operating system modules should normally be installed as part of the IPL procedure.

100.02 CONTINUED . . .

- FDRW48 INVALID REPLY--PLEASE REPLY YES OR NO**
Reason: The reply to the FDRW47 message was invalid.
Action: The FDRW47 message is reissued. Reply YES or NO.
- FDRW49 PLEASE ENTER THE FULLY QUALIFIED FDR *libname* DATA SET NAME**
Reason: The TSO allocation CLIST ABRALLOC was unable to allocate one of the FDR ISPF libraries whose name is stored in the CLIST. "libname" indicates which library was not found.
Action: If the FDR ISPF library names have changed, please enter the correct data set name for the library. You should also update the ABRALLOC CLIST with the correct names ([see Section 90](#)).
- FDRW50 TSO RESTORE FAILURE--DSN=*dsname***
Reason: The ABR Catalog Locate Exit was attempting to restore the data set "dsname" in the foreground under TSO and the restore failed. This will be accompanied by another message detailing the cause of the restore failure and either FDRW51 or FDRW52 if retry of failing TSO foreground restores is enabled.
- FDRW51 RESTORE WILL BE ATTEMPTED IN THE BACKGROUND.**
Reason: See FDRW50. The installation has enabled retry of failing foreground recalls in the background (LXFGERR=BG), so the recall will be attempted as a background task.
- FDRW52 RESTORE REQUEST WILL BE PLACED IN THE REMOTE QUEUE.**
Reason: See FDRW50. The installation has enabled retry of failing foreground recalls via the remote queue (LXFGERR=RQ), so the recall will be added to the Archive Restore remote queue for later processing.
- FDRW53 *description* TASK LIMIT REACHED -- *action***
Reason: The installation has enabled options LXMAXSTC=nnn and/or LXMAXREC=nnn (ISPF panel A.I.4.11.3) to limit the number of auto recall started tasks or the total number of auto recall tasks or both. "description" may be either RECALL STARTED or TOTAL RECALL depending on which limit was exceeded. "action" may be either
 REPLY 'WAIT' OR 'CANCEL' if the recall is for a TSO user
 WILL WAIT for all other environments
Action: The TSO user can reply with WAIT to wait for the current recalls to drop below the limit or CANCEL to abort the recall. Non-TSO recalls will automatically wait for the recalls to drop below the limit; the message is informative.
- FDRW54 CVAF I/O ERROR SEARCHING THE VTOC OF DISK VOLUME *vvvvvv***
Reason: The ABR Catalog Locate Exit issued a CVAF call to see if a data set marked as having been Archived by ABR actually exists on disk volume "vvvvvv". If the data set is on disk, the restore is not necessary. The CVAF call failed with a return code of 16 (x'10'), indicating a permanent I/O error was encountered searching the VTOC. The VTOC of that volume may be damaged.
Action: See IBM diagnosis manuals for recovery procedure. If unable to correct, call Innovation for technical assistance.
- FDRW56 SYNCHRONOUS RESTORE *status***
Reason: The ABR Catalog Locate Exit has initiated a synchronous recall on behalf of a TSO or ROSCOE user. This message reports the status of the recall. "status" may be:
 IS PROCEEDING. The synchronous restore is still in process. This message is issued at five-second intervals up to 4 times.
 TIMED OUT. The synchronous restore task failed after notifying the exit that the restore had started.
 CONVERTED TO ASYNCHRONOUS RESTORE VIA ATTENTION. The user depressed the ATTN or PA1 key. Rather than canceling the restore the WAIT is cancelled and the restore is allowed to proceed asynchronously (without impacting the terminal user).
- FDRW57 UNCATALOGED FOR DELETE--DSN=*dsname***
Reason: The ABR Catalog Locate Exit detected a catalog DELETE request for a data set Archived for auto-recall. The DELETE is failed with a return code of 8 and reason code of 42 (record not found). However, the auto-recall catalog entry is removed (uncataloged). If the delete was issued by IDCAMS or TSO you may receive:
 IDC3009I ** VSAM CATALOG RETURN CODE IS 8
 IDC0551I ** ENTRY dsname NOT DELETED
Action: None, since the catalog entry has been removed. You may disregard the IDC messages.

100.02 CONTINUED . . .

**FDRW58 DATA SET IN USE OR BEING RECALLED--
RETRY ALLOCATION IN A FEW MINUTES**

Reason: The ABR Catalog Locate Exit issued an ENQ to see if the data set listed in the FDRW70 message was available for use. The ENQ showed that the data set was allocated to another job or user. The recall is terminated.

Action: The user can retry the recall after allowing some time for the other task to free the data set.

FDRW59 ABR AUTO RECALL CANCELLED--reason

Reason: The ABR Catalog Locate Exit was entered because a catalog search (LOCATE) was issued for a data set that was Archived by ABR with the RECALL=YES option. The exit failed trying to invoke ABR to recall the data set. "reason" will be one of the following:

ATTACH OF pppppppp FAILED--RC=nnn. The exit attempted to ATTACH the program "pppppppp" but the ATTACH failed. The ATTACH return code from register 15 is printed ("nnn").

FDROPT LOAD FAILED--R1=nnnn/R15=nn. The exit attempted to LOAD FDROPT but it failed. The contents of the LOAD SVC return code registers 1 and 15 are printed, with register 1 showing the abend code that would have resulted if the FDROPT were executed as a processing program.

I.E.: // EXEC PGM=FDROPT

FDROPT VERSION nnnn NOT SUPPORTED. "nnnn" is the version number of the FDR Global option table (FDROPT) that was found. It is not compatible with the version of the exit installed.

SYSZTIOT ENQUEUE ERROR--RC=nnn. "nnn" is the return code from a failing ENQ SVC.

ATTACH NOT ATTEMPTED--SP251 IN KEY 0. The exit needed to attach a module that is linked non-reentrant. It would be loaded into subpool 251. This subpool has been acquired in protect key zero (0).

Action: If unable to correct, call Innovation for technical assistance.

FDRW60 FDROPSPF VER x.x/xx - PROCESSING ERROR - PARM-LVL level

Reason: The FDR Install dialog CLIST invoked program FDROPSPF with an incompatible parameter level, indicating that the DDNAME allocations for the FDR CLIST library and the FDR program library point to libraries at different maintenance levels. On a re-install, the problem might be that the new FDR program library is not allocated to either ISPLLIB or STEPLIB DDNAMEs, thus causing the prior level FDR programs to be loaded from the Linklist.

Action:

1. If the FDR libraries are allocated at TSO Logon time by the Logon PROC or Logon CLIST, verify that all required DDNAMEs are allocated, as documented in [Section 90](#).
2. If the FDR libraries are allocated by a local CLIST that uses the ISPF LIBDEF service, switch to use the ABRALLOC CLIST supplied by IDP, as documented in [Section 90](#).

100.02 CONTINUED . . .

FDRW69 **reason--DSN=dsname**
Reason: The ABR Catalog Locate Exit was attempting to reallocate a data set recalled by ABR at step initiation. Before the recall, the exit temporarily allocates the data set to a volume. If ABR recalls the data set to a different volume, then the exit must reallocate the data set; if it was part of a concatenation, it must re-concatenate all the data sets in the list. The attempt to reallocate or re-concatenate the data set failed for one of the following reasons:
LOCATE FAIL-RC=nn. A catalog LOCATE issued during reallocation failed with return code "nn".
DSNAME MISMATCH. The data set name returned by LOCATE is not the name the exit was expecting. This may occur if the data set has an alias.
NOT RESTORED. ABR did not successfully restore the Archived data set.
DEALLOCATE FAILED. Dynamic deallocation of the DDname assigned to the data set failed with a return code other than zero.
REALLOCATE FAILED. Dynamic re-allocation of the DDname assigned to the data set failed with a non-zero return code.
RECONCAT FAILED. Dynamic re-concatination of the DDname assigned to the data set failed with a non-zero return code.
SWA RESCAN FAILED. A scan of SWA (System Work Area) during re-allocation failed with a non-zero return code.
CVAF FAIL - nn/mm. A CVAF call was issued to verify that a data set is on disk, but it failed with return code "nn" and reason code "mm".
DEVTYPE MISMATCH. ABR recalled the data set to a different device type than the exit expected.
VOLSER NOT FOUND. ABR indicated that it recalled the data set to a specific volume serial, but the exit cannot find that volser in any online UCB.
TOO MANY DDNAMES. The data set was referenced by more than 20 DDnames in this job step. The first 20 DDnames have been reallocated.
ddname BYPASSED -- "ddname" is either JOBLIB or STEPLIB. It was not reallocated because it is currently OPEN.

Action: Except for the NOT RESTORED reason, the required data sets have been recalled by ABR. You should be able to resubmit the job for successful execution.

FDRW70 **ELIGIBLE FOR RECALL--DSN=dsname**
Reason: The ABR Catalog Locate Exit determined that a TSO user referenced a data set which has been Archived by ABR with the RECALL=YES option. This message and messages FDRW71 through FDRW79 (as required) are issued to the TSO user.

100.02 CONTINUED . . .

- FDRW71 TYPE 'END' TO BYPASS THE RESTORE OR PRESS 'ENTER' TO CONTINUE**
Reason: If option LXALTMMSG=NO, this format of message FDRW71 is issued asking the user if the recall of the Archived data set should be done. If option LXFOREST=NO, the message is not issued and the recall automatically proceeds. FDRW71 is only issued to TSO users.
Action: Press the ENTER key to recall the data set or type 'END' to terminate the recall.
- FDRW71 TYPE 'YES' TO PERMIT THE RESTORE OR PRESS 'ENTER' TO BYPASS**
Reason: If option LXALTMMSG=YES, this format of message FDRW71 is issued asking the user if the recall of the Archived data set should be done. Unlike the standard FDRW71 message above, the user must positively request the recall. If option LXFOREST=NO, the message is not issued and the recall automatically proceeds. FDRW71 is only issued to TSO users.
Action: Type "YES" to recall the data set or press the ENTER key to terminate the recall.
- FDRW71 TYPE 'UC' TO UNCATALOG THE DATA SET**
Reason: If option LXUNCAT=YES, this format of FDRW71 may also be issued, giving the user the option of uncataloging the data set instead of recalling it. FDRW71 is only issued to TSO users.
Action: Type "UC" to uncatalog the data set without recalling it. The data set will no longer be auto-recallable but it can still be restored manually. A new data set with the same name can be created.
- FDRW72 *restore-option--REPLY 'xx'***
Reason: This message is issued one or more times to show the user a list of recall techniques (FOREGROUND, BACKGROUND, REMOTE QUEUE) available to TSO users, and the reply necessary to invoke each one (FG, BG, RQ). Option LXFDREST can be set to limit the user's choices, only the recall techniques indicated will be listed. If LXFDREST has been set to a single technique, message FDRW72 is not issued and that technique will be automatically used. FDRW72 is only issued to TSO users.
- FDRW73 TYPE THE APPROPRIATE RESPONSE OR 'END' TO TERMINATE THE RESTORE**
Reason: This message follows FDRW72. FDRW73 is only issued to TSO users.
Action: Enter the 2-character response that corresponds to the recall technique selected; it must be one in the list displayed by FDRW72.
- FDRW74 ABR AUTO RECALL ABNORMALLY TERMINATED COMP CODE=00ssssuuu**
Reason: The ABR Catalog Locate Exit was entered to recall an Archived data set, but ABR abnormally terminated trying to restore the data set. The completion code contains:
 sss--is the system ABEND code
 uuu--is the user ABEND code (in hex).
 FDRW74 will be issued for batch jobs as well as online users.
Action: If unable to correct, call Innovation for technical assistance.
- FDRW75 RESTORE BYPASSED -- JOB *jjjjjjj* DATE: *yyddd* TIME: *hhmmss* DSN=*dsname***
Reason: Data set was scheduled for recall during step initiation but the recall was never attempted because the job step flushed, probably due to a JCL error.
- FDRW76 DATA SET IS SCHEDULED TO BE RESTORED TO VOLUME SERIAL NUMBER *vvvvvv***
Reason: The data set listed in the FDRW70 message is scheduled to be restored to the disk volume with serial "*vvvvvv*". If option LXNEWVOL is set to NO, the message is not issued and the recall to the volume selected by the ABR Catalog Locate Exit will automatically proceed.
- FDRW77 TYPE NEW VOLUME SERIAL NUMBER AS *xxxxxx* OR PRESS 'ENTER' TO CONTINUE**
Reason: This message is issued after message FDRW76 and asks the TSO user to either accept the target volume displayed in FDRW76 or designate a new target volume. If option LXNEWVOL is set to NO, the message is not issued and the recall to the volume selected by the ABR Catalog Locate Exit will automatically proceed.
Action: Press the ENTER key to accept the volume shown in FDRW76 or type in the volume serial of a target volume. Note that ABR or SMS may choose to restore the data set to a different volume. If option LXREISSU is set to YES, when you enter a new target volser, messages FDRW76 and FDRW77 are re-issued so that you can confirm the choice of target volume.

100.02 CONTINUED . . .

- FDRW78 RESTORE *recall type*--YOU WILL BE NOTIFIED UPON COMPLETION**
Reason: The ABR Catalog Locate Exit is recalling the data set named FDRW70 message. The user will receive additional messages indicating the success or failure of the recall. The "recall type" may be:
START COMMAND ISSUED. An asynchronous (background) recall has been started. The TSO user can continue with other functions, and can use the recalled data set when additional messages are received indicating the recall is completed.
STARTED IN FOREGROUND. A synchronous recall in the foreground has been started. The TSO user will be unable to proceed until the recall is complete.
STARTED SYNCHRONOUSLY. A synchronous recall using an external recall address space has been started. The TSO user will be unable to proceed until the recall is complete; however, the user may convert the recall to an asynchronous recall by pressing the ATTN or PA1 keys.
- FDRW79 RESTORE *reason***
Reason: The ABR Catalog Locate Exit has determined that the recall of the data set named in the FDRW70 message cannot be done for one of the following reasons:
BYPASSED - TSO RESTORE DISABLED. Auto recall of Archived data sets is disabled for TSO users (option LXDFREST is set to NO). The message is not issued if option LXNOMSG is set to YES.
BYPASSED - DEFAULT RESTORE TYPE IN ERROR. The user made an invalid reply to message FDRW73 and the default recall technique was improperly specified. The user will be prompted for another technique.
PROMPTED - DEFAULT RESTORE TYPE IS DISABLED--the default recall technique established for TSO users is disabled.
UNCATALOG FAILED - R15=ffffff R0=00000000 R1=11111111. The user replied UC to the FDRW73 message to uncatalog the data set without recalling it, but the uncatalog failed. The contents of registers 15 (return code), R0 (catalog return code) and R1 are shown.
Action: Correct if this is an error in the setup of the RECALL exit.
- FDRW80 CONFIRM REQUEST TO COMPAKT UNIT=uuuu,VOL=vvvvvv, JOB=jjjjjj**
Reason: COMPAKTOR job "jjjjjj" is about to COMPAKT the volume "vvvvvv" on device "uuuu". To suppress this message, specify CONFMESS=NO on the COMPAKT statement.
Action: Reply 'YES' if COMPAKTOR is to continue. Reply 'NO' if the COMPAKTion is to be bypassed.
- CONFIRM REQUEST TO RECOVER VOL=vvvvvv FROM FAILED COMPAKTION --
 REPLY YES OR NO**
Reason: A Fast COMPAKTion (CPK TYPE=FASTCPK) was in progress on the indicated disk volume but that COMPAKTion did not complete for some reason, such as a system crash. This Fast COMPAKTion job will automatically recover from the failure if the operator gives permission. To suppress this message and proceed with the recovery automatically, specify RECOVERY=YES on the COMPAKT statement.
Action: Reply YES to allow the recovery to proceed; reply NO to bypass the recovery at this time..
If you reply NO, the volume will not be in a usable state. Many data sets may be unusable.
You must submit another Fast COMPAKTion job and allow it to recover from the failure, or recover the volume from a backup.
- FDRW81 DETECTS UNIT=uuuu,VOL=vvvvvv IN USE nnn. REPLY 'RETRY' OR 'EOJ'**
Reason: The count of open DCBs "nnn" against volume "vvvvvv" on device "uuuu" is checked by COMPAKTOR. If the count is greater than 1, indicating that COMPAKTOR is not the only user of the volume, this message is issued. This message is also issued if the volume contains an active paging data set, or is the system residence volume. The keyword ACTMESS=NO on the COMPAKT statement will suppress this message and proceed with the COMPAKTion automatically. See the topic "COMPAKTing Active Volumes" in [Section 40](#).
Action: If you know which job is using the device, reply RETRY after it ends. If you are not sure, the MVS console command
 D U,,ALLOC,uuuu,1
 will display all jobs which have the disk allocated (the number of allocations may differ from the number of open DCBs). If necessary, reply EOJ to terminate the COMPAKTOR job. Except for the sysres and paging volumes, you can also reply IGNORE to continue the COMPAKTion; however, you must have followed the rules for "COMPAKTing Active Volumes" in [Section 40](#).

100.02 CONTINUED . . .

FDRW82	<p>COMPAKTOR <i>function status</i> VOL=vvvvvv Reason: If LOG=YES is specified on the COMPAKT statement, this message is issued to inform the operator of the progress of the COMPAKTOR job. "function" will be either "REORGANIZATION" for TYPE=FASTCPK or CPK, "SPACE RELEASE" for TYPE=RLSE or "RECOVERY" for recovery from a failed Fast COMPAKTion. "status" will be "STARTED", "SUCCESSFUL" or "***FAILED***". The "STARTED" message will remain non-deletable on the console until the "SUCCESSFUL" message is issued so that the operator will know that the operation is still in progress. The "***FAILED***" is also non-deletable to force the operator to take action.</p> <p>Action: If the status is **FAILED**, check the output from the COMPAKTOR job. The volume should not be used until any problems are resolved (See Section 40.18).</p>
FDRW88	<p>text Reason: This message is issued by the FDRSPZAP program. The text will indicate various FDRSPZAP information or errors.</p>
FDRW89	<p>FDR -- TRIAL VERSION FROM INNOVATION DATA PROCESSING EXPIRES IN <i>nnn</i> DAYS Reason: This is a trial version of the FDR system. The number of days the trial will remain active is displayed. When there are 10 or fewer days before the trial is due to expire this message will become non-deletable.</p> <p>Action: When the trial expires it will not be usable. If you have licensed a production version of FDR, you should install it in place of the trial version. Call Innovation if you need assistance.</p>
FDRW90	<p>DSN=<i>dsname</i> NOT CATALOGED. REASON CODE=<i>reason</i> Reason: This message is issued by module FDRPRE00 (the ABR DADSM Pre-exit) when it fails to catalog a data set being scratched or renamed in the ABR Scratch catalog. The exit tried to catalog a name consisting of '#' followed by the user's dsname, in order to record the ABR backup information.</p> <p>REASON CODE EXPLANATION</p> <ul style="list-style-type: none"> A An I/O error occurred reading the VTOC of the volume containing the data set. B A catalog LOCATE on the ABR Scratch catalog failed. This message also contains the VSAM catalog return and reason codes. Refer to message IDC3009I in the IBM System Messages manual for the meaning of the codes. Common causes for this message are: <ol style="list-style-type: none"> 1) The ABR Scratch catalog may be full. Run the Scratch catalog maintenance utility (FDRABRCM PURGE SCRATCH, Section 50) to cleanup obsolete entries. 2) A security system may have disallowed the operation. All users must be allowed to catalog names starting with '#.' C A catalog update on the ABR Scratch catalog failed. See Reason B. D The data set name being processed is all blanks. E The catalog LOCATE for the ABR Scratch catalog alias failed. The data set name displayed in the message contains the ABR Scratch catalog alias name; it may have not been properly specified in the master catalog as an alias of the Scratch catalog. F An old (non-ICF) VSAM catalog was used as a STEPCAT. G Error creating a GDG base in the ABR Scratch catalog for recording a particular generation data set being scratched. See Reason B. K The DADSM scratch or rename workarea was not found. L RESERVE/DEQ failed on SYSVTOC. M An I/O error was reported by DADSM rename. N See message FDRW94. O The data set name being scratched or renamed was 43 or 44 characters long, so that the exit was unable to add the "#." Scratch prefix without exceeding the 44-character limit. Such data sets cannot be recorded in the Scratch catalog. Set option PXERRLEN to NO to suppress this message (see Section 90). P VTOC TTR conversion error. Q The value indicating the displacement in the DSCB for the recording of old backup information is different between the FDR Global Option Table and the ABR Model DSCB of the volume involved. R Unable to record DELETE VVR. S The current generation number in ABR Model DSCB is zero. T FDRPRE00 GETMAINs an area which is freed by FDRPOST0 (the DADSM post processing exit). The area was not freed. U Unable to ENQ on SYSZTIOT in a multi-tasking address space. V GQSCAN failure. W FDRPOST0 (DADSM post-processing exit) was not invoked during a scratch. <p>Action: This message may be accompanied by an IEC331I message. If necessary, contact Innovation for assistance.</p>

100.02 CONTINUED . . .

- FDRW92 MAXIMUM NUMBER OF ERRORS EXCEEDED LOG TERMINATED**
Reason: FDRPRE00, the ABR DADSM pre-processing exit, has issued 100 FDRW90 error messages since the last time the exit was installed (usually the last IPL). No further messages will be issued until you do a re-IPL or reinstall the exit.
- FDRW93 FDR DATA SET NOT FOUND EXIT ABNORMALLY TERMINATED COMP CODE= cccccccc**
Reason: The ABR Data Set Not Found (DSNF) exit was entered for a data set whose name is contained in the preceding FDR316 message. ABR failed to recall the data set from the ARCHIVE backup. Either the ABR return code (in hex) or the abend code are displayed in this message. If the first two characters of cccccccc are '0F', then the ATTACH SVC for FDRABR failed, and the remaining digits contains the hexadecimal return code from the ATTACH SVC.
Action: Contact Innovation for assistance.
- FDRW94 FDRPRE00 ESTAE RECOVERY ENTERED FOR ABEND Ssss Uuuuu location**
Reason: This message is issued by FDRPRE00 when an unexpected condition occurs while trying to record a scratched or renamed data set in the ABR Scratch catalog. It documents the system (sss) or user (uuuu) ABEND that occurred. "location" indicates where the error occurred, either an offset within FDRPRE00 or "IN NON-IDP EXIT" if the error occurred in another DADSM exit from another vendor (or locally-written).
Action: Call Innovation for assistance.
- FDRW99 IF function IS CANCELLED resource MAY BE CORRUPTED - REPLY options**
Reason: An MVS CANCEL command was issued against an executing FDR program. Some FDR programs contain CANCEL protection since a resource may be corrupted if the CANCEL is allowed to take place. The programs are FDRREORG, FDRCPK, and FDRARCH. "function" may be: FDRREORG, FDRCPK, or the FDRARCH function being executed. "resource" may be DATA SETS, VOLUMES, VOL=vvvvvv or ARCHIVE. "options" will contain the reply options available for this message.
Action: Depending on the available reply options displayed, the operator can reply:
S Stop at the end of the current operation. FDR will complete processing of the data sets or volumes currently being processed and no new work will be started.
I Ignore the CANCEL command. The FDR job will continue normally.
C Accept the CANCEL. The FDR job will be cancelled and the indicated resource may be corrupted. **This should only be done in emergency situations or when the FDR job appears to be hung (not doing any I/O or consuming any CPU time).**
WARNING: replying C to COMPAKTOR will leave the volume in an unusable state.
Another form of the FDRW99 message will be issued to confirm your response. Note that FDRREORG also accepts the console P (STOP) command which is the equivalent of the S reply ([see Section 30.06](#)).

100.03 MESSAGES FROM FDR, DSF AND ABR (FDRnnn)

- FDR001** *program/function -- VER v.r/mmt -- INNOVATION DATA PROCESSING DATE=yyyy.ddd PAGE nnn*
Reason: This is the FDR page heading, containing the name of program or FDR function generating the message and the version level of FDR. "v.r" indicates the version and release (e.g., 5.3), "mm" is a 2-digit number indicating the maintenance level and "t" will be "P" for a production version or "T" for a trial.
- FDR002** *function SUCCESSFULLY COMPLETED VOL=vvvvvv NVOL=nnnnnn*
Reason: The disk with volume serial vvvvvv was successfully dumped, restored, copied, moved or simulated. On a full volume restore, vvvvvv is the volume serial of the disk from which the dumps were created and nnnnnn is the serial number of the receiving volume before the restore.
- FDR003** **NONSTANDARD RECORD ZERO -- cccchhhhrkklll -- function CONTINUING**
Reason: FDR detected a non-standard record zero (R0) in a track on the volume. The count field of the R0 is printed in hex. A standard IBM record zero format has an rkklll field of 00000008 -- that is, record number zero, key length zero, and data length eight. FDR will not RESTORE a track that had a non-standard record zero during the backup.
Action: Contact Innovation for assistance.
- FDR006** *function OPERATION CANCELLED BY OPERATOR*
Reason: An operator replied 'NO' to an FDR/DSF request for approval to restore or copy a disk. A U0801 Abend follows.
- FDR007** *startend TIME OF function -- hh.mm.ss - UNIT=disktype, IN=inputdd ,OUTPUT=outdd1 outdd2*
Reason: Documents the time that FDR, DSF, or ABR began or ended a dump, restore, or copy. "startend" will be STARTING or ENDING. "function" will indicate the type of function (e.g., DATA SET DUMP). "inputdd" is the input DD name. "outdd1" is the output DD name and "disktype" identifies the input or output disk device type (in the case of emulated disk, such as RAMAC, the emulated disk type, such as 3390, will be shown. If a duplicate backup was requested, it is shown by DD name "outdd2". On a data set restore or copy/move, since data sets may be output to multiple devices at once, the disk type and output DD names are not shown (output volumes are shown in message FDR311).
- FDR008** **OPEN ERROR OR NO DD STATEMENT DD=ddname - function BYPASSED**
Reason: 1. A required disk or tape DD statement specified by "ddname" was missing.
2. DDname SYSPRINx is missing.
3. An error occurred while OPENing the specified "ddname".
Action: The disk in error will be bypassed. See the joblog for possible IBM OPEN error messages.
- FDR010** **POSSIBLE SPACE ERRORS IN VTOC -- ENTIRE PACK WILL BE DUMPED**
Reason: FDR detected an I/O error or logical error reading the VTOC. A mini-dump is printed detailing the error. FDR will continue processing, dumping the entire disk volume.
Action: Contact Innovation for assistance.
- FDR012** **NOT AN FDR TAPE DSN=dsname**
Reason: 1. The input for a restore ("dsname") is not an FDR created backup.
2. The JCL specifies the wrong serial numbers, or specifies the serials in the wrong order. FDR must always read the backup from the beginning; it cannot start in the middle of a multi-volume backup. If you specified the volume serials of the backup tape in JCL, check the order of the serials.
3. The FDR backup file may have been written over.
Action: "dsname" is the tape data set name. A mini dump is printed displaying the first tape block read. A U0205 Abend is issued.
Hint: the first block of a valid FDR backup is a 24 byte record containing the text beginning "THAT'S ALL FOLK". If anything else is displayed, see the reasons above.

100.03 CONTINUED . . .

FDR014 TAPE BACKUP IS INCOMPATIBLE WITH RESTORE DEVICE

Reason: An attempt was made to do a full-volume restore or copy to an unlike device type (e.g.: 3380 to 3390). Use FDRDSF, FDRCOPY, or FDRABR data set restore to restore to a device with a different device geometry. It can also occur when attempting to do a full-volume restore to a smaller disk of the same type (e.g., 3390-3 to 3390-2).

Action: The FDR header read is printed in a mini dump. A U0104 Abend follows. If necessary, contact Innovation for assistance. If you want to force FDR to restore to a smaller disk of the same type, specify PROT=NONE on the RESTORE or COPY statement. The FDR014 message will still be printed, but the RESTORE or COPY will be attempted. You will get I/O errors if any tracks above the capacity of the smaller disk must be restored. If all of the data on the larger disk will fit on the smaller disk, use COMPAKTOR to restore and relocate the data.

FDR019 RACF FACILITY PROTECTION FOR BYPASS FAILED FOR *resource*

Reason: On a system with SMS active, the operands BYPASSACS or BYPASSSMS were specified. To use these the user must be authorized to the RACF resources
STGADMIN.ADR.RESTORE.BYPASSACS (for restore)
or STGADMIN.ADR.COPY.BYPASSACS (for copies)
in class FACILITY or the equivalent in other security systems.

Action: A control card error is issued.

FDR020 RACF VOLUME PROTECTION FAILED ON VOL=*vvvvvv*

Reason: Security checking was enabled in the FDR Global Option Table. A full volume dump or restore was requested but a security call for class DASDVOL and volume "*vvvvvv*" failed. The user does not have sufficient authority to perform this operation. This message is produced only for full-volume restores and absolute track operations; full-volume backups and data set operations will do class DATASET checks against individual data sets if the DASDVOL check fails.

Action: A U0801 Abend is issued.

FDR021 RACF DATA SET PROTECTION FAILED ON DSN=*dsname*

Reason: Security checking was enabled in the FDR Global Option Table. A security call for class DATASET and data set name "*dsname*" failed. The user did not have sufficient authority to access this data set.

Action: On a full volume backup operation, the backup is terminated. On a data set operation, this data set is bypassed.

FDR022 RESTORE TAPE IS A DSF FORMATTED TAPE

Reason: A full volume restore was requested from a backup created by a FDRDSF or FDRABR data set backup. This backup may not contain an image of the entire volume.

Action: Normally FDR will terminate the RESTORE with a U0205 ABEND. If TAPE=DSF was coded on the RESTORE control statement, FDR will continue the RESTORE. Warning messages will be issued if the volume LABEL or VTOC were not found on the tape.

FDR023 DUMMY VTOC READ -- ENTIRE PACK WILL BE DUMPED

Reason: FDR found that the VTOC on this volume started and ended on cylinder zero head zero. This is the format used on volumes initialized by VM; that dummy VTOC is not valid.

Action: FDR will dump all tracks on the volume.

100.03 CONTINUED . . .

- FDR024 INVALID EXTENT DESCRIPTOR [REASON=reason DSN=dsname]**
Reason: FDR read a DSCB (Format 1 or 3) which contained an invalid extent description. If a minidump is printed, the first five bytes of the DSCB printed after the registers in the mini-dump are the cylinder, head, and record number (CCHHR) of the DSCB in error. If REASON= is printed, it is one of the following:
1 - ENDING CYLINDER TOO LARGE
2 - STARTING TRACK TOO LARGE
3 - ENDING CCHH BEFORE BEGIN
4 - ENDING TRACK TOO LARGE
 Note: Reasons 2 and 4 indicate that the starting or ending track number of an extent is higher than the highest track on a cylinder, i.e. higher than 14 on 3380 or 3390; not that the start or end of an extent is after the end of the pack.
Action: Depending on the operation, FDR may continue the backup and dump all tracks on the volume, or may ABEND with a U0110.
 Scratch the bad data set.
 If an absolute track statement was specified (SELECT FROM/TO), the user specified a cylinder address greater than maximum cylinder number or a track number greater than the tracks per cylinder. Correct the SELECT statement.
 If necessary, contact Innovation for assistance.
- FDR030 INVALID PARM FIELD PASSED TO FDR**
Reason: The user specified a PARM parameter to program FDR but the value was invalid. [See Section 10.03](#) for valid parameters.
Action: A U0405 ABEND follows. Correct PARM field and re-submit job.
- FDR031 DD=ddname -- NUMBER OF CHARACTERS IS NOT 5 -- STATEMENT IGNORED**
Reason: FDR encountered a DD statement starting with the characters 'DISK' followed by more than 1 character, e.g., DISK12. For FDR and FDRDSF, DISK can be followed by exactly one alphanumeric or national character. In FDRABR, you can place up to 4 characters after DISK, so this message is not produced.
Action: The DD statement is ignored and FDR will attempt to process the remaining 'DISK' DD statements.
- FDR032 NO VALID DISK(X) DD (OR MOUNT) STATEMENTS WERE FOUND**
Reason: FDR did not find any DD statement starting with the characters 'DISK'.
Action: FDR did not dump any disk volumes.
- FDR033 CONTROL OPTION CONFLICTS WITH PARM FIELD OR ANOTHER OPTION -- JOB TERMINATED**
Reason: 1. If you specified both a PARM= for FDR and control statements in SYSIN, one indicates a dump and the other indicates a restore.
 2. The RESTORE statement contains CPYVOLID=YES but PARM=N is in the JCL.
 3. A COPY statement was provided, but PARM=x is in the JCL.
 4. The input contains more than one DUMP, RESTORE, or PRINT statement.
 5. MAXTASKS=n was specified on a COPY operation.
- FDR034 MAXIMUM NUMBER OF DISK(x) DD STATEMENT EXCEEDED -- JOB TERMINATED**
Reason: More than 39 DISKx DD statements were provided in the JCL. FDR processed the first 39 disk volumes.
Action: A U0502 Abend is issued.
- FDR035 NO DSNAME CONTROL STATEMENTS FOUND -- FDR TERMINATED**
Reason: The control statements specified to FDRDSF did not specify that any useful work was to be done. User did not specify that any data sets were to be dumped. At least one SELECT statement must be specified or DSN=ALL specified on the DUMP control statement.
Action: Correct control statements and re-submit job.
- FDR036 UNSUPPORTED DISK DEVICE TYPE ON DDNAME=ddname**
Reason: FDR encountered a disk device or a backup of a disk device which is not supported. It is possible that a back level release of FDR is being used.
Action: Contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR040 VOLUME IN USE -- nnn OPENED DCBS**
Reason: FDR on a full volume dump or restore found the volume currently being used by other jobs. The number of opened DCBs (minus one for FDR) is printed. This is the count in the UCB at the beginning of the dump or restore. This is only a warning message, no error code is issued.
Action: If you do not want this message printed, the operand ACTMESS=NO on the DUMP statement will suppress the message.
- FDR042 RESTORE FROM BACKUP OF VOL=vvvvvv
 CREATED ON DATE=yyyy.ddd TIME=hh.mm.ss**
Reason: Documents the creation time and date of the backup from which a restore is being done, as well as the original volume serial (vvvvvv) of the disk that was dumped.
- FDR090 function - VER v.r/mmt - INNOVATION DATA PROCESSING DATE/TIME=yyyy.ddd/hh.mm.ss
 PAGE nnn**
Reason: General page heading for the FDR mini-dump processor listing the version, level, date, time and page. See message FDR001 for details.
- FDR091 REGS**
Reason: Data formatted by the mini-dump processor, printing from 1 to 16 general registers and, optionally, related storage. All register oriented data is identified by this message number.
- FDR092 type LEN=nnnn(hhhh) LOC=lllll**
Reason: Control block formatted by the mini-dump processor. 'type' is the name of the control block, 'nnnn' is the length in decimal, 'hhhh' is the length in hex, and 'lllll' is the storage location in hex.
- FDR093 CCWS LEN=nnnn(hhhh)**
Reason: CCWs formatted by the mini-dump processor.
- FDR094 reason -- CCWS NOT PRINTED**
Reason: An FDR related program requested a mini-dump with CCW formatting. However, for the reason printed, CCW printing was terminated.
Action: Contact Innovation for further problem determination.

100.03 CONTINUED . . .

- FDR099 NEAR RELATIVE DATA LOCATION *nnn* - error description**
Reason: The FDR common parsing routine encountered an error in parsing a control statement. The approximate location of the error was position *nnn*, counting the first position as 000. The failing statement is normally printed immediately above.
Action: Correct error described in 'error description' and re-submit job.
- FDR101 *program/function* -- VER *v.r/mmt* -- INNOVATION DATA PROCESSING DATE=*yyyy.ddd* PAGE *nnn***
Reason: This is the DSF page heading. See message FDR001 for details.
- FDR102 PROGRAM IS NOT APF AUTHORIZED**
Reason: Most FDR programs must be executed as an APF-authorized program in order to execute correctly. This program detected that it was not executing authorized. This most often occurs when you are testing a new version of FDR with a STEPLIB which has not been authorized.
Action: 1) authorize the library from which you executing the FDR program. You can update a library list in SYS1.PARMLIB and you may be able to input a console command which will authorize the library.
 2) move all FDR programs to an authorized library.
 Contact Innovation if you need assistance.
- FDR106 DATA SET NOT ON VOL=*vvvvvv* DSN=*dsname***
Reason: FDR did not find the Format 1 DSCB for data set "*dsname*" in the VTOC of volume "*vvvvvv*", but it expected it to be there.
Action: There are a variety of unusual conditions which can produce this message. Call Innovation for support.
- FDR107 *function* SUCCESSFULLY COMPLETED VOL=*vvvvvv***
Reason: The indicated function was successfully completed on the disk volume "*vvvvvv*".
- FDR108 WARNING -- LABEL OR TRACK IN VTOC RESTORED X'*cccchhhh*'**
Reason: A DSF absolute track restore or copy (SELECT FROM/TO) has overlaid the volume label track (cylinder 0 head 0) or a track in the VTOC. The cylinder (*cccc*) and head (*hhhh*) of the overlaid track is printed in hex.
Action: If you intended to restore this track, the message can be ignored. If not, the volume may not be usable until you recover the overlaid track from a backup.
- FDR109 TAPE NOT FORMATTED FOR DATA SET RESTORE**
Reason: A DSF backup created by absolute track dump (SELECT FROM/TO) was input to a data set restore.
Action: A U0202 Abend follows. You can only do absolute track restores from such backups.
- FDR111 VOL=*vvvvvv* DSN=*dsname* reason**
NUMBER OF TRACKS NECESSARY=*tttt*
EXTENT CYL=*cccc* TRK=*ttt*,CYL=*cccc* TRK=*tttt*
Reason: This message is printed for data set restores when a pre-allocated output data set is too small or has other errors. It will print the first two forms of the message shown above, and then repeat the last form for all extents of the input data set (up to 128).
 "reason" can be:
HAS INSUFFICIENT SPACE *nnnnn* TRACKS - the data set is too small to hold all of the tracks to be restored. "*nnnnn*" is the size of the pre-allocated dataset. The next FDR111 message shows the number of tracks to be restored "*tttt*". This can also occur for data sets allocated by FDR if the CYL= or TRK= operands specified a size too small for the active tracks of the data set, or when DATA=ALL is specified with the RLSE or %FREE= operands.
IS UNMOVABLE - the input data set is flagged as unmovable, which usually means that it contains location-dependant data and cannot be restored to different track addresses. The additional FDR111 messages will show the extents that the data set occupied on the original volume; it must be restored to the same track addresses.
IS TOO LARGE - an ICF VSAM index component was larger than the index on the backup. It must be the same size.
Action: The data set will not be restored. Scratch the pre-allocated data set. In most cases, simply re-run the restore and FDR will allocate the data set properly. In some cases you may need to pre-allocate the data set with the allocations shown (for example, FDR cannot allocate an unmovable data set with more than one extent). If you get this message for a data set allocated by FDR (not pre-allocated), contact Innovation for assistance.
- FDR120 INVALID DEVICE TYPE -- *function* TERMINATED**
Reason: The disk device type is not one supported by this version of FDR.
Action: A U0103 Abend follows. Contact Innovation for assistance.

100.03 CONTINUED . . .

**FDR121 ALTERNATE ERROR REASON=n ON TRACK CYL/HEAD=ccccchhhh
ASSIGNED TO PRIME CYL/HEAD=ccccchhhh PRIME R0=ccccchhhh**

Reason: FDR detected an error or unusual condition with the alternate track assignments on the disk volume being processed. The message displays the cylinder and head number of the failing alternate track in hex. ASSIGNED TO PRIME is printed if FDR can determine the prime track; it will display the cylinder and head of the prime track and the count field of record 0 (R0) on the prime track (the latter should point to the alternate track. This message is only produced when you are taking a backup to a sequential disk data set (such as Archiving to disk) or when you specify BUFNO=nn.

The reason code "n" can be:

- 1 I/O error occurred reading this alternate track. If the IOB in the mini-dump starts with '46008000', the volume has no alternate tracks (this occurs with some disk subsystems and with VM mini-volumes). As long as there are no alternate tracks assigned on the volume, this message can be ignored.
- 2 This alternate track is flagged as a defective alternate. FDR will bypass it. No corrective action is required.
- 3 The alternate track is assigned to a prime track. The alternate home address does not have the flag byte set indicating that this is an alternate track.
- 4 I/O error occurred reading the prime track assigned to this alternate.
- 5 The prime track's home address flag byte did not indicate that this is a defective track.
- 6 The prime track's record zero (R0) did not point back to this alternate track. (PRIME R0=) contains data read. If the PRIME R0 points to another alternate track, this track has had an alternate assigned to it multiple times.

Note: In many cases (especially reason 1), this message is informative and does not indicate an error unless other disk errors occur; in this case, the message does not set a non-zero return code or cause an ABEND.

Action: FDR will continue, bypassing this alternate. If the prime data track is to be dumped, FDR may issue additional error messages. FDR121 REASON=1 can usually be ignored. Contact Innovation for assistance.

**FDR122n statistics
Reason:**

FDR displays status information at the successful termination of a backup or restore subtask. A second FDR122 is printed displaying the values under each header. The character "n" after the FDR122 indicates the format of the backup and will be N-new format or S-split format. These headers will appear above the statistics values (some values are displayed only for backups, not for restores):

BYTES	Number of bytes of data dumped to tape or read from tape (for DSF, all bytes read even if not restored.)
DSK TRK	Number of disk tracks accessed to perform this function.
T BLKS	Number of tape blocks written or read. On COPY this is an internal number of blocks passed from the dump task to the restore task.
RESTART	Internal field
STIMERS	Internal number of STIMER commands issued.
ERRS	Using BUFNO=MAX, this is the number of tracks in error.
ACT DSK	Actual number of disk tracks dumped or copied. For FDRCOPY, it will be the number of tracks read, which is not necessarily the number of tracks written.
LOW	If not BUFNO=MAX, this is the lowest number of STIMERS issued on one cylinder. If BUFNO=MAX, this is the number of retries on tracks with data checks. Unless a FDR129 error message is displayed, the retries were successful.
HGH	If not BUFNO=MAX, this is the highest number of STIMERS issued on one cylinder. If BUFNO=MAX, this is the number of count fields encountered with a CCHH which did not match the track it resides on; this is normal on a VM mini-disk but may indicate errors on other volumes.
DEXCP	The number of disk EXCPs issued by this subtask to dump this volume. This number may be less than the number of EXCPs recorded by SMF. For FDRCOPY, this is the total of EXCPs on the input and output disks.
NUMDS	Number of data sets processed in this backup. If an ABR incremental backup, this value includes the backup of the ABR model DSCB.
COMP BYTES	The number of bytes to which the data was compressed by FDR software compression (COMPRESS=). This field is not displayed if software compression was not requested (FDR cannot display the results of hardware compression, e.g., IDRC)). It is displayed only for backups.

100.03 CONTINUED . . .

FDR122
(continued)

OPERATION STATISTICS FOR type VOLUME...volser
Reason:

This form of the FDR122 message may appear as the header for a table of statistics about the certain FDR operations that were performed on the volume indicated; currently the previous form is used for backup, restore and print operations and this form is used for most other operations (such as SNAP). The table may include some or all of the data below. For non-backup/restore operations (such as SNAP), many of the statistics printed may be zero or very small; this is normal.

CYLINDERS ON VOLUME - total number of data cylinders on volume being processed

DATASETS PROCESSED - number of data sets selected.

BYTES READ FROM DASD - total number of bytes actually backed up from the volume

BYTES WRITTEN TO DASD - total number of bytes actually restored to the volume

BYTES ON BACKUP - total number of bytes actually written to or read from the backup. For uncompressed backups, this will be slightly higher than the DASD bytes because of control data. For compressed backups (COMPRESS= on the DUMP statement), this is the bytes after compression.

COMPRESSION SAVINGS - percentage of the backup file size saved by FDR compression (COMPRESS= on the DUMP statement). This is calculated as the "Bytes on Backup" divided by the bytes that would have been written to the backup without compression (not displayed). This will be zero unless COMPRESS=ALL was specified; compression savings from IDRC tape hardware compression is not displayed

DASD TRACKS BACKED UP|WRITTEN - number of data tracks read from or written to the volume.

BACKUP BLOCKS WRITTEN|READ - number of physical blocks written to or read from the backup file.

DASD EXCPS - number of write/read I/O requests issued to the volume.

BACKUP FILE EXCPS - number of write/read I/O requests issued to the backup file.

CPU TIME (SECONDS) - the CPU (TCB) time required to process this volume, in seconds and thousandths of a second.

ELAPSED TIME (MINUTES) - the actual time, in minutes and tenths, required to process this volume.

BACKUP|RESTORE TIME (EXCLUDING MOUNTS) - the actual time, in minutes and tenths, required to process this volume, excluding the time spent waiting for input or output tape volumes to be mounted (at OPEN or EOVS).

BACKUP COPY n ON type DSN=dsname VOL=volumes - documents the location of the backup file created. "n" is the copy number (1 or 2) and "type" is DISK or TAPE.

RESTORE FROM DSN=dsname - documents the name of the backup file being restored.

100.03 CONTINUED . . .

FDR123 COUNT RECORD ERROR ON TRACK X'cccchhh' REASON=n

Reason: FDR detected an error on the count fields of the cylinder (cccc in hex) indicated. Reason codes 1 through 6 can only occur if BUFNO=MAX is not in effect (BUFNO=MAX is the default except when outputting to disk or dumping with FORMAT=SPLIT). The track number (hhh in hex) printed may not be accurate. An I/O error may also occur during the actual dump of data from the track. If a cylinder has more than one count error, only one FDR123 message will be issued per track. The description of "DISK I/O ERRORS" in [Section 80](#) may be helpful in understanding the error.

The reason code "n" may be:

- 1 The number of count fields read on this cylinder exceeded the maximum number FDR expects for this device. The remaining count fields will be bypassed.
- 2 I/O error occurred reading the count fields. The track number in the message is probably not correct. THE ENTIRE CYLINDER WILL BE BYPASSED. The IOB in the mini-dump details the error.
- 3 The number of bytes on this cylinder and head exceeds the maximum set by FDR. If it does not exceed the FDR maximum buffer size FDR will process this track. **WARNING: This track may cause an I/O error when reading the data or the track may not be restorable.** If the number of bytes exceeds the buffer size, the entire cylinder will be bypassed.
- 4 Count field is out of sequence. The count field encountered has a head number lower than the previous count field. FDR will process this record as if it was the next sequential record. This track may not be correctly dumped.
- 5 The length of the count fields on this track developed into a negative number. FDR will terminate immediately with a U0300 Abend.
- 6 The track number on one or more count fields was higher than the maximum number of tracks per cylinder for this device (usually 15). FDR will bypass the records from the point of error to the end of the cylinder. The track number in the message has no significance.
- 7 Count field has an incorrect cylinder and track number using BUFNO=MAX. Since this may be a normal condition (such as on VM volumes), this reason is printed only if you specify PRINT=ALL on the DUMP statement.
- 8 The data read from this track appears to exceed the capacity of the track. The data read will still be dumped. This track may cause an error when restored.

Action: An FDR mini-dump will be printed. The unlabeled block of storage at the end of the mini-dump contains the count fields read. Each count is 8 bytes (2 words) long and consists of the CC (cylinder), HH (head), R (record), K (key length) and DD (data length). The count field causing the error is usually the first one on the third line of counts. Contact Innovation for assistance.

FDR124 FORMAT 4 ERROR REASON=n

Reason: FDR detected a problem with the Format 4 DSCB on a volume. The Format 4 DSCB describes the VTOC itself. The VTOC may be improperly formatted.

The reason code "n" may be:

- 1 More than one Format 4 DSCB was found, or the first DSCB in the VTOC was not a Format 4. This may be due to an IBM problem.
- 2 FDR found that the Format 4 contained zero free DSCB's plus the VTOC high water mark was set to the first Format 5 DSCB. This is the condition in which CPK would leave the volume if an ABEND or system crash occurred during the restore phase. This volume may be partially restored. **This is a very serious problem. The volume should be placed offline until the error is corrected.** See [Section 40](#) for a procedure to recover the volume. If you must dump the volume, specify PROT=NONE on the DUMP statement.
- 3 Label track is invalid. The label track (cylinder 0 track 0) does not contain a properly formatted volume label.

Action: A U0602 Abend will be issued. List the VTOC in hex with

```
//LIST EXEC PGM=IEHLIST
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=3390,VOL=SER=vvvvvv,DISP=OLD
LISTVTOC VOL=3390=vvvvvv,DUMP
```

and contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR125 I/O ERROR READING THE VTOC X'ccccchhh' VOL=vvvvvv
TRACK WILL BE BYPASSED**
Reason: FDR detected an I/O error reading the VTOC on the cylinder and track indicated in hex. An FDR mini-dump is printed. The DSCBs stored on this track of the VTOC will not be processed. This error may also be caused by an improperly formatted VTOC track. CPK cannot restore from this backup.
Action: Correct the VTOC or contact Innovation for assistance.
- FDR126 ERROR PROCESSING FORMAT 2 OR 3**
Reason: FDR attempted to read the Format 2 or 3 DSCB associated with a Format 1 DSCB. The VTOC location pointed to by the Format 1 DSCB does not contain a DSCB of the proper type. The DSCB chain for this data set may be broken.
Action: Correct or scratch this data set or list the VTOC in hex with

```
//LIST EXEC PGM=IEHLIST
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=3390,VOL=SER=vvvvvvv,DISP=OLD
LISTVTOC VOL=3390=vvvvvvv,DUMP
```

and contact Innovation for assistance.
- FDR127 DISK TRACK CROSS-CHECK ERROR CCHH=ccccchhh - DUMP CONTINUING**
Reason: FDR detected that the record length fields developed for a tape block were inconsistent with the count fields of the data dumped. The cylinder and track (in hex) of the failing track is displayed.
Action: Contact Innovation for assistance.
- FDR128 INVALID RECORD ZERO ON TRACK X'ccccchhh' DATA X'ccccchhh'**
Reason: The track specified has an invalid record zero (R0). The count field of R0 should always contain the ID of the track it resides on (except for tracks with assigned alternate tracks). The cylinder and track (in hex) of the failing track is displayed, along with the count field of R0.
Action: FDR will bypass this track on the disk and continue processing the volume. Contact Innovation for assistance.
- FDR129 I/O ERROR ON DISK PACK--LAST HOME ADDRESS READ X'ccccchhh'**
Reason: An I/O error or logical error was detected by FDR on this volume. The cylinder and track (in hex) printed may not be accurate. An IBM IOS000I message may also have been printed on the Job log.
Action: An FDR mini-dump is printed detailing the I/O error ([see section 100.01](#)). FDR will attempt to continue processing from this point, bypassing the track in error. An FDR150 message may also be issued. If the error is a data check or other hardware problem, the volume must be corrected. See Disk I/O Errors in [Section 80](#) for help in understanding the error. Contact Innovation for technical assistance.
- FDR130 CYL=cccc HEAD xxx...x WAS function**
Reason: Internal message on each cylinder and head dumped or restored. An X appears for each track accessed within the cylinder. "cccc" is the cylinder number in decimal. "function" is either **ACCESSED** (read) or **RESTORED, COPIED** or **MOVED**.
- FDR131 YOU ARE NOT LICENSED FOR FUNCTION - function**
Reason: You attempted to execute a FDR function, but that function is not included in your current FDR license. "function" indicates the failing function, such as REORG (FDRREORG), INSTANTBKP (FDR InstantBackup), DCT=YES (HSDM support) and others.
Action: If you believe you are licensed for this function, contact Innovation technical support for assistance. Check to be sure that you are executing FDR from the library containing the latest copy of FDR installed at your site. If you would like to test an unlicensed product, contact your Innovation sales representative to request a trial.

100.03 CONTINUED . . .

- FDR138 UNABLE TO FUNCTION DSN=*dsname* REASON=*reason***
Reason: A **RESTORE, COPY** or **MOVE** did not complete successfully. "reason" is a reason code plus explanatory text as shown below. Some of the reasons below occur only during logical restore (to an unlike device or when reblocking) but others can occur on any data set restore.
 The reason code "reason" may be:
- 4 CONTAINS TRACK OVERFLOW RECORDS** -- data sets formatted with track overflow records (supported only on 3350 and earlier disks) cannot be restored logically)
 - 7 UNSUPPORTED DSORG** -- these data sets are not supported for logical restore.
 BDAM ABSOLUTE KEY RANGE OR SPANNED
 NON-ICF VSAM VSAM MULTI-VOLUME
 ISAM NOT SUPPORTED -- see member ISAM in the ICL library
 - 9 BLOCK SIZE LARGER THAN TRK SIZE** -- on a logical restore, a data block larger than the output track size was found. Use BLKF= if the data set is DSORG=PS
 - 9 TOTAL BLOCKS LARGER THAN TRK SIZE** -- on a logical restore where track image mode was forced (see message FDR139), the data from one track of the input data set will not fit on a track of the output data set. Message FDR151 will be printed to identify the output track. This message may also occur if an input data set was identified as having multiple extents not in ascending order on the backup or input disk, indicating that track image mode was attempted but failed, probably because of very small blocks in the input data set (for a PDS from a 3380 being restored to a 3390, this will occur if the PDS has a directory with over 45 blocks). See "Action".
 - A CATALOG/VVDS/PAGE TO UNLIKE** - logical restore cannot restore catalogs, VVDSs or page/swap data sets
 - A SYS1.LOGREC OR VTOCIX TO UNLIKE** -- logical restore cannot restore VTOC indexes or LOGREC data sets.
 - B DATA SET IS UNMOVABLE** - logical restore cannot restore data sets flagged as unmovable
 - F SPACE IS TOO SMALL FOR DATA SET** -- during a logical restore, the allocated space was too small for the reformatted data. If FDR allocated the data set, this may be due to very small blocks in the data set. If possible, FDR will attempt to extend the data set to a secondary extent. The CYL= or TRK= operands may be used to increase the size of individual data sets
 - L I/O ERROR WRITING TO DATA SET** - another message will document the error
 - Q REBLOCK ERROR *reason*** - a logical restore error occurring while trying to reblock a data set (the BLKF= operand of RESTORE). "reason" briefly describes the error
 - R PDS DIRECTORY *reason*** - a logical restore error occurred while processing the directory of a PDS (DSORG=PO). "reason" briefly describes the error.
 - T TAPE BLOCK DEBLOCK ERROR** - an FDR data block was improperly formatted. The backup may have been corrupted.
 - U END OF FILE NOT FOUND ON INPUT** -- during a logical restore, an expected EOF record was not found within the dumped tracks of the data set; the data set may be usable.
 - V VSAM ERROR *reason*** - a logical restore error occurred while processing an ICF VSAM cluster. "reason" briefly describes the error.
 - X DATASET WAS EXCLUDED**
 - Y TRACKS OUT OF ORDER FROM INPUT** -- on a logical restore, the tracks of the original non-VSAM data set were in multiple extents and not in ascending order on the backup or input disk. See "Action".
- Action:** Correct the error, if possible, and rerun. If this was a logical restore or copy/move to an unlike device, it may be possible to do a physical restore/copy/move to a like device (same device type). If the error was due to multiple extents not in ascending order on a non-VSAM data set (REASON=Y or REASON=9), it may be possible restore or copy/move to a like device to reallocate the data set in a single extent then copy/move it to the unlike device (or use COMPAKTOR on the input disk to combine the extents of the data set). Read member UNLIKE in the ICL (Installation Control Library) for more details. Contact Innovation if you need assistance.

100.03 CONTINUED . . .

- FDR139 TRACKIMAGE FORCED DSN=*dsname* REASON=7
BDAM VARIABLE OR UNDEFINED**
Reason: A logical restore of a BDAM (DSORG=DA) data set with variable or undefined format will be done in track-image mode, restoring tracks exactly as they were dumped.
Action: None. This is a warning only unless other errors occur.
- WARNING DSN=*dsname* REASON=7 CLUSTER NOT LOADED**
Reason: In this form of message FDR139, an attempt was made to perform a logical restore of an empty ICF VSAM cluster. FDR will only allocate and catalog this cluster.
Action: None. This is a warning only.
- WARNING DSN=*dsname* REASON=W MISSING TTR *ttttr* FROM MEMBER=*member***
WARNING DSN=*dsname* REASON=W MISSING NOTE LIST TTR *ttttr*
Reason: In this form of message FDR139, during a logical restore or REORG of a PDS, a TTR pointer (relative track "ttt" and record number "rr", both in hex) in the PDS directory or in a note list within a member did not point to any record in the original PDS. This is an error which existed in the PDS before FDR restored or reorganized the PDS.
Action: The TTR will not be changed. It will still be invalid after the restore/reorg. All other members will be processed. You should investigate the indicated member; it should probably be deleted.
- WARNING DSN=*dsname* REASON=U END OF FILE NOT FOUND ON INPUT**
Reason: In this form of message FDR139, during a logical restore of a PS or DA file, no EOF was found in the data set. However, all tracks were restored so the data set should be valid.
- WARNING DSN=*dsname* DATA=ALL IGNORED DUE TO DUMP OPTIONS**
Reason: In this form of message FDR139, during a logical restore DATA=ALL was specified, but the dump was not done with DATA=ALL.
Action: DATA=ALL is ignored. The data set should be usable.
- REORG BYPASSED *reason* DSN=*dsname***
Reason: In this form of message FDR139, a FDRCOPY REORG was requested for a data set, but it could not be done for the "reason" indicated:
 PDSE PDSEs are not supported for REORG
 UNMOVABLE the data set's DSORG was POU
 MODEL the data set has no extents
 NOT A PDS the data set was not partitioned
- FDR142 ERROR USING THE READ MULTIPLE CKD COMMAND (5E) -
NEW PERFORMANCE OPTION TERMINATED**
Reason: FDR detected an error using the 5E (RMCKD) command. This control unit may not support the command.
Action: FDR will revert to the previous performance option. Contact Innovation for assistance.
- FDR149 tracedata**
 Displays data from an internal trace. The FDR operands to invoke various traces will be provided by Innovation when the data is required to diagnose a problem you have reported. The message is also sometimes used when I/O errors occur.

100.03 CONTINUED . . .

- FDR150 DISK TRACK X'cccchhh' BYPASSED--ERROR WHEN DUMPED**
Reason: FDR detected an I/O error or logical error when this track was dumped. Another message may have been issued detailing the error. Since the data was not properly dumped, FDR will not restore the track on the restore function.
Action: The data set associated with this track may be incomplete. On a data set restore, message FDR155 will indicate the dsname. If message FDR155 does not appear, the track was not needed for this restore. If you need assistance, contact Innovation.
- FDR151 I/O ERROR DSN=dsname
I/O ERROR WRITING DISK TRACK X'cccchhh' VOL=vvvvv
TRACK BYPASSED**
Reason: FDR detected an I/O error writing the disk track at cylinder cccc track hhhh (both are in hex). For a data set restore only, the first FDR151 message will identify the data set which contains that track.
Action: FDR will continue the restore operation. This track will be bypassed. See Disk I/O Errors in [Section 80](#) for help in understanding the error. Contact Innovation for technical assistance.
Note: If this is a logical restore and you also received an FDR138 REASON=9, then refer to the instructions in the Action for the FDR138 Message.
- FDR152 ICF VSAM ERROR REASON=x DSN=cluster/component VOL=vvvvv**
Reason: An I/O error or logical error exists within the SYS1.VVDS data set or in the VVDS information for the cluster or component indicated, on volume "vvvvv" (this error may also occur for non-VSAM data sets on SMS-managed volumes). On a restore, FDR compares the characteristics of the cluster on the backup file with those of the cluster being restored to on disk; if some required characteristic does not match, this message details that mismatch. The VVR/NVR from the backup and the disk will be printed for certain types of errors.
The reason code "n" may be:
- 1 ICF VSAM clusters exist on the volume, but FDR failed to find the SYS1.VVDS data set on the output disk or on the backup file. The VVDS may not exist or FDR encountered an error trying to find it. The IOB details the error.
 - 2 The number of extents in the SYS1.VVDS data set was zero.
 - 3 The SYS1.VVDS contains more than 3 extents. FDR had an I/O error or logical error finding a Format 3 DSCB. The IOB details the error.
 - 4 An I/O error or logical error was detected when FDR read or wrote entries within the SYS1.VVDS data set, or the VVDS contains no active records (even though there are ICF clusters on the volume). For an I/O error, the IOB details the error.
 - 5 The length of a VVR/NVR entry within the VVDS is negative.
 - 6 The length of all of the VVR/NVR entries within a block of the VVDS exceeded the length in use within the block.
 - 7 A VVDS record indicates that the length of its in-use data exceeds 4096.
 - 8 An error occurred when DSF attempted to find one of the associated components for this cluster. The component may be in the VVDS but not in the VTOC. The IOB details the error.
 - 9 FDR is unable to store all of the VSAM component and cluster names in its internal table. A large number of VSAM clusters must exist on the volume. Specify the keyword ICFCORE= on the DUMP TYPE= Statement or permanently increase the ICFCORE value in the FDR Global Option Table.
 - A A cluster specifies a catalog name which was not found in the catalog entries stored in the first record of the VVDS.
 - B FDR could not find the volume information/allocation cell (type 23) in the VVR record, or an individual cell had an invalid length.
 - C More than one SYS1.VVDS data set found on the volume. The operating system will use the VVDS which contains the current disk volume serial number in the name. More than one VVDS can be created if a volume with a VVDS is renamed and a VSAM cluster is allocated using the volume with the new serial number. If the VVDS with the incorrect volume serial does not have any current clusters cataloged within it, use SUPERZAP to turn off the PASSWORD indication in the DSCB and scratch the data set using IEHPROGM. A backup will complete but no ICF VSAM clusters can be restored from it and non-VSAM SMS data sets may not be restorable. A restore will issue a U0634 ABEND, so no ICF VSAM or SMS data set restores can be done to this volume until the condition is corrected.
 - D The volume serial in the data set name of the VVDS (SYS1.VVDS.Vvolser) does not match the volume serial of the disk. This can occur if a volume was copied or restored to a new disk volume serial but the VVDS was not renamed. This is only a warning; FDR will backup the volume properly, using that VVDS. However, VSAM and SMS data sets on the volume will not be usable until the condition is corrected.
 - E FDR did a GETMAIN for above the line storage for a VSAM table but the GETMAIN failed. Try increasing the REGION= parameter to more than 32M.

100.03 CONTINUED . . .

- F** FDR did not find both a VVR/NVR and a DSCB for a given data set; one was missing. Also occurs when FDR was unable to calculate the free space in a cluster.
- G** The imbedded index VVR (Type Q) was not found.
- H** The size of an extent was not a multiple of the CA size; or the high RBA of an extent exceeded the maximum. On a logical restore, the original cluster had an imbedded index (the IMBED attribute) and the high-level index component had more than one extent.
- I** The control interval size (CISIZE) on the backup cluster does not match the cluster on disk. The displacement is X'16' (See Note 1).
- J** The number of CIs per CA on the backup cluster does not match the cluster on disk. The displacement is X'0E' (See Note 1).
- K** The space allocation unit (TRACK vs CYL) on the backup cluster does not match the cluster on disk. The displacement is X'10' (See Note 1). This error is also issued following a FDR160 message; see FDR160.
- L** The physical blocksize or number of blocks per track on the backup cluster does not match the cluster on disk. The displacement is X'11' or X'15' (see Note 1).
- M** The number of tracks per CA on the backup cluster does not match the cluster on disk (Check the secondary allocation value specified). The displacement is X'17' (see Note 1).
- N** The VSAM cluster is being restored to a different type of cluster, e.g., KSDS to an ESDS. Flags at 2 displacements are checked (see Note 1). At displacement X'03' it checks:
 X'80' - Key Sequence (KSDS)X'20' - Imbedded Index (IMBED)
 X'10' - Replication (REPLICATE)X'04' - Key Range
 X'02' - Relative Record (RRDS)
 At displacement X'2A' it checks:
 X'04' - Linear (LDS)X'01' - Variable RRDS (VRRDS)
 If none of these flags are on, the cluster is an ESDS (entry sequenced) cluster. There may be other flags on in those bytes; they don't matter. But all of the flags named above must match (on or off) between the input and output cluster.
- O** The length of the VVR/NVR exceeds the length of the cells within the VVR/NVR.
- R** SMS Attributes of the backup cluster and the cluster on disk do not match. FDR checks for Extended Format (EF) and for "over 4GB addressing". The displacement is either X'03' or X'1D'.
- S** The backup cluster contains a key range while the disk does not or the key lengths are not the same. The displacement is X'24' (see Note 1).
- T** A VVR entry in the VVDS contains a cell type other than a 21, 23 or 60.
- U** The type of VVR entry on the backup does not match the disk (Z record being restored to Q). The user may be attempting to restore a multi-volume component to the wrong sequence component.
- V** The total length of a VVR/NVR record does not match the sum of the lengths of the cells within it, on the output disk.
- W** The total length of a VVR/NVR record does not match the sum of the lengths of the cells within it, on the backup file.
- X** The number of extents specified in the VVR record does not match the DSCB entry in the VTOC for a component.
- Y** A base cluster is being restored to an alternate index or vice versa.
- Z** Internal Error -- The length of the FDR backup file control block for VSAM was less than the sum of the lengths of the VVRs/NVRs within it.

Action:

- 1) During a dump, the cluster/component indicated (or all clusters if SYS1.VVDS.Vvvvvv is the name printed) may not be restorable. A FDR mini-dump may be printed to help diagnose the error. .
- 2) During a restore, the cluster or component indicated will not be restored (it is possible for one component to be restored without error while another fails, but the cluster will not be usable). For many reason codes, an FDR mini-dump will be printed displaying the VVR from the backup and the VVR from the disk (see Note 1).
- 3) For I/O errors, an FDR mini-dump will be printed showing the I/O results.
- 4) Errors resulting from a comparison of the characteristics of the cluster being restored compared to cluster on disk are actually not detected until the restore is complete; the data tracks have been restored but the cluster is probably not usable. See Note 1.
- 5) If restoring to a pre-allocated cluster, it was not defined with the characteristics that FDR requires to restore; in most cases, simply delete the cluster and let FDR define it properly. Some errors can occur even if FDR allocates the cluster. Contact Innovation for assistance if the error cannot be corrected. See VSAM Special Considerations ([Section 80](#)).

100.03 CONTINUED . . .

Note 1: For the indicated reason codes, FDR will print a mini-dump displaying the VVR (VSAM Volume Record) from the backup file and from the cluster on disk. In the registers at the top of the dump, register 14 points to a cell within the disk VVR and register 15 points to a cell within the backup VVR. Those registers plus the hex displacements shown above for the appropriate reason codes will point to the fields which did not compare.

FDR153 VSAM SCRATCH ERROR COMP=cccc CODE=code CLUSTER=cluster

Reason: FDR attempted to scratch an ICF VSAM cluster but the CATALOG DELETE SVC failed. "cccc" and "code" match the return code and reason code documented for IBM message IDC3009I except that "cccc" values over 9000 are internal FDR errors.
Some commonly occurring IBM codes are:

COMP CODE

0008 0042 The cluster to be scratched was not cataloged. FDR can only scratch cataloged clusters
0076 0000 User attempted to scratch a multi-volume cluster which was not aliased to a user catalog and a STEPCAT was not present for the catalog. This data set may be partially deleted.
0076 0008 User attempted to scratch a cluster with an alternate index which has the NOUPGRADE attribute on another volume.
0084 0000 The VSAM cluster is date protected. Specify VEXPD=NONE on the DUMP Statement to scratch this cluster.

COMP values over 9000 are from FDR:

COMP

9004 Dynamic allocation failed for the specified cluster. "code" is the decimal equivalent of the dynamic allocation error code; it also appears in hex in register 14 within the mini-dump. Dynamic allocation error codes are documented in the ISPF online tutorial appendix and in various IBM manuals. In particular, a dynamic allocation error code of decimal 5896 or X'1708' indicates that the input cluster was not found in the catalog. If the cluster is not aliased to a user catalog in the master catalog, you may need supply to a STEPCAT DD statement for the proper catalog.
9008 User attempted to move a multi-volume VSAM cluster. The components on this volume were copied to the output volume, but were not scratched. If all of the associated components have been dumped or copied, the user can use SUPERSCRATCH or IDCAMS to scratch the cluster.
9012 Internal Error. The TIOT search failed to find the disk DD name to which the cluster was dynamically allocated.
9016 FDR cannot scratch an ICF VSAM cluster on a non-MVS system.
9020 A KSDS with alternate indexes has been archived but it was not scratched because ALTINDEX=NO was specified.
9024 Cluster cannot be scratched because it is cataloged to a different volume or because an alternate index with the NOUPGRADE attribute is on a volume by itself.
9028 Cluster cannot be scratched because it is cataloged as non-VSAM or cataloged to more than 255 volumes.

Action: Check the error code and correct if necessary. An FDR mini-dump will be printed displaying the CATALOG DELETE parameter list and the catalog name. If needed, contact Innovation for assistance.

FDR154 ARCHIVE SCRATCH BYPASSED SOME VOLUMES NOT PROCESSED CLUSTER=clustername

Reason: A multi-volume VSAM cluster was backed up from some, but not all, of the volumes on which it resides. The cluster will not be deleted; it is still usable. However, some parts of it have been recorded in the Archive Control File. ABR will terminate with a non-zero return code unless SELTERR=NO is specified on the DUMP statement.

Action: You must insure that all volumes on which the multi-volume cluster are processed in the same Archive step. One way to do so is to select the cluster with SELECT CATDSN=cluster. This message may occur if you are trying to select clusters with ADAYS= or ADATE=; this does not work for multi-volume VSAM since the last reference date is stored only in the Format 1 DSCB of the base data component and only on its first volume. See VSAM Special Considerations in [Section 80](#).

100.03 CONTINUED . . .

- FDR155** **nnnnn TRACKS MISSING ON function OF DSN=dsname**
Reason: Data set **RESTORE, COPY** or **MOVE** detected that a number of expected tracks were not found during the restore of the data set "dsname". "nnnnn" is the number of tracks that were not found. The remainder of the data set was restored, but it may be unusable. The cause may be:
1. The restore specified DATA=ALL, but the backup used DATA=USED
 2. The restore used DATA=ALL or DATA=USED but the backup was created with DATA=NONE
 3. The restore used LBPZERO=INVALID, but the backup used LBPZERO=VALID
 4. Backup file errors (message FDR200) were reported during the backup or restore
 5. A disk error occurred during the backup which required tracks or cylinders to not be recorded on the backup file (e.g., message FDR123 or FDR129). Message FDR150 will also be printed for such tracks during the restore
 6. Blocks were lost from the backup file due to unreported tape drive errors
 7. The restore subtask abended before completing the restore (message FDR997 is also issued).
- Action:** If the cause cannot be determined, contact Innovation for assistance.
- FDR156** **ALLOCATE FAILED FOR nnnnn quan COMP=X'code-reason' VOL=vvvvvv DSN=dsname**
ALLOCATE FAILED FOR CATALOG COMP=comp CODE=reason VOL=vvvvvv DSN=dsname
Reason: FDR attempted to allocate a non-VSAM data set "dsname" on volume "vvvvvv" and the attempt failed (if FAILED FOR CATALOG is displayed, the allocation was successful, but a later attempt to catalog the data set failed; this can also occur for VSAM). "nnnnn" is the size of the data set which failed allocation in tracks or cylinders, "quan" will be TRK or CYL. If the allocation was attempted on multiple volumes (because of the NVOL= operand or the ABR RESTORE ALLOCATION LIST), it failed on all of them, but "vvvvvv" is the first volume on which it was tried, and "comp" is the error code from that volume. On a system with SMS active, if the data set is SMS-managed, "vvvvvv" is the volume that FDR selected for the data set and may not be the same as the volume on which SMS tried to allocate it. "comp" may be return codes from the IBM ALLOCATE SVC (SVC 32), from the IBM CATALOG SVC (SVC 26), from the IBM SMS interface, or internal diagnostic codes from FDR itself.
- If FAILED FOR CATALOG is displayed, "comp" is the return code and "reason" is the reason code from the CATALOG SVC, in decimal, as documented for the IBM message IDC3009I. Some of the common catalog codes shown under message FDR157 also apply to FDR156. Otherwise, if "comp" is a value less than X'1000', it is from the ALLOCATE SVC, in hex. "reason" may be a 8-digit hex IBM diagnostic code. For OS/390 and ESA see the IBM manual DFSMSdfp Diagnosis Reference.
- Common ALLOCATE codes include:
- 0004** Duplicate Data set name in VTOC; data set already exists on the volume. This may occur for multi-volume data sets if the piece currently on the volume has a different volume sequence number than the piece being restored.
 - 0008** VTOC or VTOC index (VTOCIX) is full
 - 000C** I/O error in VTOC or VTOC index (VTOCIX)
 - 0010** Requested absolute tracks not available; an unmovable data set or one which was allocated with ABSTR cannot be allocated because the required tracks are not free
 - 0014** Requested space quantity not available; DADSM was unable to find sufficient free space (in up to 5 extents) to allocate the data set or the specific space required for an unmovable data set was not free.
 - 00AC** No security authorization to create data set. You must have ALTER authority to the data set under RACF or the equivalent in other security systems
 - 00B4** IGGPRE00 installation exit rejected the request. A local exit or DASD Control software product has disallowed allocation on the volume that FDR selected
 - 00C0** SMS failed the allocation request. The most common cause is a mismatch between the SMS flags in the VTOC of the output volume and that volume's status in the current SMS configuration; common reason codes are:
 - 04160053** - VTOC indicates INITIAL status (in process of conversion to SMS)
 - 04160054** - VTOC indicates NON-SMS, but the volume is in a SMS storage group.
 - 04160055** - VTOC indicates SMS, but the volume is not in a SMS storage group
 You can use FDRCONVT ([Section 70.30](#)) to convert the volume to SMS or non-SMS, or you can remove the volume from its storage group in the current SMS configuration. "comp" values greater than 1000 are from FDR, in hex, as follows:
 - 2354** The name of the data set being restored does not meet IBM's data set naming conventions. If NEWNAME=, NEWGROUP= or NEWINDEX= was specified, the error is in the new name; otherwise it is the original name of the data set which is invalid. You can bypass this check by specifying NODSNCHK on the RESTORE statement; the data set will be allocated but the catalog will probably fail

100.03 CONTINUED . . .

- FDxx** On a system with SMS active, the IBM subsystem interface (IEFJSREQ) failed with return code "xx"
- FExx** On a system with SMS active, the IBM SMS interface failed the allocation request with return code "xx", (usually 08). "reason" will be an SMS error code, in decimal. You can find the description by looking up IBM message IGDxxxxx, where xxxxx is the reason code
- FF04** FDR cannot allocate an ISAM file. See member ISAM in ICL Library
- FF08** FDR cannot allocate a non-ICF VSAM file
- FF0C** FDR cannot allocate an unmovable file with more than 1 extent
- FF10** FDR cannot allocate a ICF VSAM file on a non-MVS system
- FF18** On a COPY/MOVE of a data set protected by a discrete RACF profile, an attempt to issue RACDEF to define a profile for the output data set failed. "reason" will be "0000xyy" where "xx" is the RACF return code and "yy" the RACF reason code ("xx" of 04 indicates that the profile already exists, 08 that the user is not authorized to create the profile)
- FF50** GETMAIN failure during allocation. Increase the region
- FF51** Allocation parameter list error. Contact Innovation
- FF52** FDR did not select a target output volume. This can occur if the data set's original volume or the volume you specified as NVOL= is not online. Specify a new NVOL= value or update the ABR Restore Allocation List to direct the data set to an online volume.
- FF53** NVR is not record type N. Contact Innovation
- FF54** NVR cell type not 22. Contact Innovation
- FF55** NVR type 24 cell length invalid. Contact Innovation
- FF56** NVS SMS subcell not type 24. Contact Innovation
- FF57** DSCB address not passed. Contact Innovation
- FF58** The selected target disk volume was SMS-managed, but no storage class was assigned. It may be that your SMS ACS routines are assigning a null storage class to this data set. To restore as SMS, specify STORCLAS=; if it still fails, use BYPASSACS or update your ACS routines. To restore as non-SMS, specify NVOL= to direct it to a non-SMS volume.
- FF59** Volsers in CPL cannot be located, or over 10 volumes passed. Contact Innovation
- Action:** Determine the cause and take appropriate action. It may be possible to circumvent the problem by pre-allocating the output data set. If needed, contact Innovation for assistance.

100.03 CONTINUED . . .

**FDR157 VSAM ALLOCATION ERROR COMP=comp CODE=reason
VOL=vvvvvv CLUSTER=clustername**

Reason: FDR attempted to allocate an ICF VSAM cluster, using the IBM CATALOG SVC, but the attempt failed. "comp" is the return code and "reason" is the reason code from CATALOG, in decimal, or is a value generated by FDR itself. If the allocation was attempted on multiple volumes (because of the NVOL= operand or the ABR RESTORE ALLOCATION LIST), it failed on all of them, but "vvvvvv" is the first volume on which it was tried, and the codes are from that volume. On a system with SMS active, if the cluster is SMS-managed, "vvvvvv" is the volume that FDR selected for the cluster and may not be the same as the volume on which SMS tried to allocate it. Codes from the IBM CATALOG SVC can be found documented under IBM message IDC3009I. Common codes include:

CODE REASON

0004	00010	The ICF catalog selected by FDR to catalog this cluster does not exist (see the description of the ICFCAT= operand). If you dumped this cluster on one system and are restoring it on another, this may mean that a catalog with the name of the original cluster's catalog does not exist; you may need to specify ICFCAT=ALIAS to catalog it into the proper catalog.
0004	00120	The ICF catalog selected by FDR to catalog this cluster is not defined in the master catalog. See 0004-0010 above.
0008	00008	Name already in the catalog. A NEWGROUP= or NEWINDEX= value may have resulted in renaming the output cluster or some of its components to the same name.
0008	00038	Name already in the catalog. Either the cluster or component name already exists, or the original cluster is still cataloged. The VRECAT operand can be used on the RESTORE statement to delete the duplicate names.
0056	00006	Insufficient security authorization to define the cluster. You must have ALTER authority to the cluster under RACF or the equivalent in other security systems.
0068	00008	The VTOC or Indexed VTOC is full.
0068	00020	Insufficient space available to define the cluster.
If "comp" is a value above 9000, the error is generated by FDR itself (COMP values from 9052 through 9080 may occur if the cluster being restored had an imbedded index, and at the time of the backup the index component was in multiple extents):		
9004		The VVR entry for the data component on the backup contains a cell type other than a 21, 23 or 60.
9012		FDR cannot allocate a multi-volume component.
9016		The VVR entry for the index component on the backup contains a cell type other than a 21, 23 or 60.
9020		FDR is attempting to allocate a KSDS cluster, but could not locate the index component on the backup. This may be a multi-volume cluster.
9024		SELECT specified a component name instead of cluster name, or a component of the requested cluster was missing from the VVDS at the time of the backup, or ICF=IGNORE or DFEF=IGNORE was specified during the dump.
9028		FDR could not find the type 23 cell in the VVR.
9032		FDR cannot allocate a VSAM cluster which contains an alternate index to a NEWNAME. Use NEWGROUP or NEWINDEX instead of NEWNAME.
9036		FDR cannot allocate a key range VSAM cluster using NEWNAME=. Use NEWGROUP= or NEWINDEX= instead of NEWNAME=.
9040		The volume serial number in the SYS1.VVDS data set does not reflect the volume it is on. If FDR were to allocate an ICF VSAM cluster under these conditions, the operating system would create a new VVDS data set.
9044		The name of the cluster or one of its components being restored does not meet IBM's data set naming standards. If NEWNAME=, NEWGROUP= or NEWINDEX= was specified, the error is in the new name; otherwise it is the original name of the data set which is invalid. This check can be bypassed by specifying the operand NODSNCHK (not documented elsewhere) on the RESTORE/COPY/MOVE command.
9048		Error occurred applying the NEWINDEX= value.

100.03 CONTINUED . . .

9052	A LOCATE SVC failed. CODE shows the return code from LOCATE. The reason code is not shown.
9056	Unable to find the PCCB for a user catalog.
9060	Dynamic allocation failed. CODE shows the return code in register 15 from SVC 99. The dynamic allocation return code (DARC) from the SVC 99 parameter list is not shown.
9064	VSAM OPEN failed. Message IEC161I may appear on the job log. CODE shows the return code from VSAM OPEN, as documented under IBM message IEC161I. For example, code 0040 indicates a security violation.
9068	VSAM PUT failed. CODE shows the return code from VSAM PUT.
9072	AMB Extension was not found.
9076	Internal error in reconstructing the imbedded index; RBA of index extent is not correct.
9080	VSAM EOV failed. Message IEC070I may appear on the job log. CODE shows the return code from VSAM EOV, as documented under IBM message IEC161I. For example, code 0104 or 0209 indicates that the volume that you were restoring to did not have enough space to contain the data set. It should be possible to restore the cluster to another volume with more space available.
9084	Error from RACF, for a cluster with a discrete profile. FDR successfully allocated the cluster, and then tried to define a new discrete profile, using the profile of the original cluster as a MODEL. The CODE value is the RACF return code, and register 0 within the mini-dump contains the RACF reason code. A CODE of 4 indicates that the profile already exists, 8 indicates the user is not authorized to create the profile.
9088	A KSDS with an excessive number of key ranges cannot be defined.
9092	A PAGE/SWAP cluster cannot be copied or moved to the same name (NEWNAME=, NEWGROUP=, or NEWINDEX= must be used).
9093	CISIZE of input cluster is invalid for logical restore.
9094	Physical blocksize of input cluster is invalid for logical restore.
9095	CASIZE of input cluster will not fit in a cylinder of the output disk for logical restore. For example, a cluster allocated in cylinders on a 3390 cannot be restored to a 3380.
9096	Allocation of a multi-volume VSAM cluster failed. This can also occur for a KSDS with the IMBED option or keyranges. "reason" indicates the specific error:
00004	-- GETMAIN/FREEMAIN error
00008	-- Volume list from LOCATE is invalid
00012	-- Non-VSAM catalog entry invalid
00016	-- Unable to obtain LOCAL lock
00020	-- Type 23/21 cell not found in VVR
00024	-- Cluster on more than 255 volumes
00028	-- Number of extents does not match Type 60 cell. Probably due to a previous incomplete restore of this cluster. See <i>VSAM Special Considerations</i> in Section 80 for details on recovering from this error.
00032	-- UCB not found for output disk
00036	-- Imbedded index extents do not match
00040	-- Length error on cluster/component name
00044	-- VVR length invalid
00048	-- NEWNAME= not supported (use NEWG/NEWI)
00052	-- Catalog entry type not non-VSAM. Multi-volume VSAM clusters which have not been completely restored will be cataloged as non-VSAM (the first volser will be ####Vx) but the cluster name is currently cataloged as something else (usually VSAM). This may be due to a previous incomplete restore of this cluster. See "Action" below.
00056	-- LOCATE error. This can occur if the original catalog of the cluster is not available on the system on which the restore was done. Specify ICFCAT=ALIAS to restore into the aliased catalog on the current system. Also see "Action" below.
00060	-- Cluster has alternate index
00064	-- VRECAT failed
00068	-- Cluster has over 123 extents on a volume
00072	-- Number of extents on volume not correct. Probably due to free space fragmentation on the target volume, causing DADSM to allocate our request in multiple extents.
00076	-- Unlike device restore not supported

100.03 CONTINUED . . .

- 00088 -- Catalog entry type not non-VSAM, and VRECAT was specified, but the piece of the cluster on the current source volume is guaranteed candidate space and does not warrant deleting the cataloged cluster. The parts of the cluster that contain data may have been restored already, or may be restored subsequently.
- 01xxx -- OBTAIN error (xxx is OBTAIN error code)
- 02xxx -- DYNAMIC ALLOCATION error (xxx is allocation error code)
- 03xxx -- EXTENT error (xxx is EXTEND error code)
- 03129 -- Insufficient space for component on this volume
- 9100** The selected output disk volume was SMS-managed, but no storage class was assigned. Specify STORCLAS= or select a non-SMS volume, and resubmit.
- 91xx** On a system with SMS active, the IBM SMS interface failed the allocation request with return code "xx" (usually 08). "reason" will be an SMS error code, in decimal. You can find the description by looking up IBM message IGDxxxxx, (where xxxxx is the reason code)
- 92xx** On a system with SMS active, the IBM subsystem interface (IEFJSREQ) failed with return code "xx".
- 9450** GETMAIN failure during allocation. Increase the region size.
- 9451** Allocation parameter list error. Contact Innovation.
- 9452** UCB address not passed or invalid. Can occur if the data set is not SMS-managed (no storage class assigned), but FDR has not selected a volume on which to allocate it. Specify NVOL= and rerun.
- 9458** BYPASSSMS was specified and the selected output disk volume was SMS-managed, but no storage class was assigned. Specify STORCLAS= or select a non-SMS volume, and resubmit.
- 9459** Volsers in CPL cannot be located, or over 10 volumes passed. Contact Innovation.

Action: An FDR mini-dump is printed displaying the CATALOG DEFINE parameter list and related data areas. Determine the cause and take appropriate action. For errors involving names already in the catalog, the VRECAT operand on the RESTORE statement may be used to DELETE the duplicate names from the catalog (but should be used cautiously since it may DELETE other clusters). If needed, contact Innovation for assistance.

When restoring a multi-volume VSAM cluster, the cluster will not be usable until all pieces of the cluster have been restored from all volumes on which it resided. If not all the pieces have been restored, the cluster will be cataloged as non-VSAM with "####Vx" as the first volsr. If possible, determine which pieces were not restored and restore them. But if you want to delete the partially restored cluster and start over, you must perform these steps:

Delete the non-VSAM catalog entry with

```
//DELNVSM EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
DELETE clustername NONVSAM NOSCRATCH
```

Delete each data and index component of the cluster from each disk to which it was restored with

```
//DELVVR EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//DD1 DD UNIT=SYSALLDA,VOL=SER=volsr,DISP=OLD
DELETE componentname VVR FILE(DD1) CAT(catalogname)
```

More details on this are found in [Section 80.13](#).

100.03 CONTINUED . . .

- FDR158 DATA SET ENQ FAILED DSN=*dsname***
DATA SET ENQ FAILED BUT HFS QUIESCED DSN=*dsname*
DATA SET ENQ FAILED BUT HFS QUIESCE FILE NOT KNOWN TO SYSTEM
DATA SET ENQ FAILED DSN=*dsname* hferror
Reason: The data set enqueue option (DSNENQ=) was in effect. This data set could not be enqueued because some other job or user currently has the data set enqueued under the SYSDSN resource.
 For HFS data sets, if HFS=QUIESCE was specified and the ENQ is not successful, FDR will try to do a HFS quiesce function, to prevent access to the file system in the HFS data set during the backup. If that is successful the second form of the FDR158 message will be displayed and the data set will be processed normally; this message does **not** indicate an error.
 If the HFS file was **not** known to the system on which FDR was running, the third form of the FDR158 message is printed, this also does **not** indicate an error.
 However, HFS errors will cause the fourth form of the FDR158 message to be printed with "hsferror":
AND HFS QUIESCE FAILED cccc-rrrrrrr – the quiesce function failed with return code "cccc" and reason code "rrrrrrr" in hex. It may be that the HFS data set was busy for more than 5 minutes. A reason code ending in 00FC indicates that the userid of the FDR job was not authorized to use Unix Systems Services (does not have a UID assigned); a UID is required.
AND HFS CURRENTLY QUIESCED TO ANOTHER JOB - the quiesce function failed because another job or user has already quiesced the HFS data set. Another backup of the HFS data set may be in progress.
AND HFS QUIESCE NOT AUTHORIZED - the job was not authorized to quiesce the data set. This may indicate that the userid of the FDR job was not defined to your security system.
AND HFS QUIESCE INSUFFICIENT STORAGE - a GETMAIN for storage to hold a table of quiesced HFS data sets failed. Increase the region and rerun.
AND HFS UNQUIESCE FAILED – at the end of the operation, the HFS unquiesce function failed. You may need to manually unquiesce the file system.
Action: In the first two cases the data set will be treated as if the SYSDSN ENQ failed.
 FDR BACKUP -- the data set will be dumped.
 COMPAKTOR -- the data set will not be moved during the COMPAKTion.
 DATASET DUMP -- the data set will be dumped unless ENQERR=BYPASS is specified.
 DATASET RESTORE -- the data set will not be restored.
 FDRCOPY -- the data set will not be copied/moved.
 ARCHIVE or SUPERSCRATCH -- the data set will be bypassed unless ENQERR=PROCESS is specified.
 ABR SNAP or SPLIT – the point-in-time backup of the volume will **not be created**. Either remove DSNENQ=USE or add ENQERR=NO to create the backup.
 An enqueue failure will cause an ABEND or non-zero condition code will be issued at the end of the FDR step (except for COMPAKTOR, ARCHIVE or SUPERSCRATCH). Specifying ENQERR=NO will suppress the ABEND or non-zero condition code caused by an ENQ failure.
- .FDR159 UNABLE TO *function* REASON=*x* DSN=*dsname***
Reason: FDR is unable to **RESTORE, COPY or MOVE** the named data set for the reason specified. The data set will be bypassed. The reason codes "x" are:
 1 An ICF VSAM catalog or SYS1.VVDS data set cannot be restored to a newname and cannot be copied or moved. To move an ICF catalog, back it up, delete the catalog, and restore it to its new volume.
 2 The SYS1.VVDS.Vvvvvvv or SYS1.VTOCIX.Vvvvvvv data set was selected without PROT=NONE specified on the RESTORE TYPE=DSF statement. It will not be restored; add PROT=NONE if you really want to restore it. It is usually a mistake to restore the VVDS or VTOCIX.
 3 A SELECT with ALLDSN or DSN=filter has selected one of these data sets:
 SYS1.VVDS.Vvvvvvv (VVDS)
 SYS1.VTOCIX.Vvvvvvv (VTOC index)
 FDRABR.Vvvvvvv (ABR model DSCB)
 Archive Control File (for restore from Archive Backup)
 Application Control File (for restore from Application Backup)
 FDR will bypass the restore of the data set. These data sets are read and updated by FDR during a restore, so it is almost always a mistake to try to restore them as data sets. SELECT ALLDSN may cause FDR159 REASON=3 for the VVDS even though it was not backed up; it can be ignored. This message is only a warning message; it will not cause an ABEND or non-zero condition code to be issued. If the VTOC index is damaged, the best way to recover is usually to rebuild it by using the BUILDIX command of program ICKDSF. To rebuild the ABR

100.03 CONTINUED . . .

- model DSCB use the REMODEL command of program FDRABRM ([Section 50](#)).
- Action:** Correct if this is an error. The message can be ignored if you did not intend to restore these data sets, but if the message is annoying, add an EXCLUDE DSN=dsname statement before the SELECT to suppress it.
- FDR160 UNABLE TO RESTORE -- INDEX HAS INCORRECT EXTENT SIZES -- CLUSTER=clustername SCRATCH THE CLUSTER AND TRY THE RESTORE AGAIN LETTING FDR ALLOCATE IT FOR YOU. MESSAGE FDR152 REASON=K FOLLOWS.**
- Reason:** The VSAM cluster that was backed up, or the input VSAM cluster for the copy or move, had an imbedded index (the IMBED option), and had an index component that was in multiple extents. In order to do the restore or copy or move, the output cluster must have the same arrangement of extents as the original cluster. However, the output cluster was found to exist on disk before the restore or copy or move, with a different arrangement of extents. There is no practical way for the user to DEFINE the output cluster with the correct arrangement of extents.
- Action:** Delete the cluster that you pre-defined, and let FDR do the allocation. FDR will automatically rebuild the cluster with the original arrangement of extents.
- FDR161 XMS ERROR REASON=x**
- Reason:** FDR was unable to find module IGG0CLX0 in the Catalog Address Space (CAS) using Cross Memory Services (XMS). The reason codes document various internal errors.
- Action:** Contact Innovation for assistance.
- FDR162 SMS ERROR MESSAGE FOR DSN=dsname message text**
- Reason:** The allocation of an SMS-managed data set failed. The diagnostic and informational messages generated by SMS during this allocation are printed.
- Action:** Examine the messages to determine the cause of the allocation failure.
- FDR164 CANDIDATE-SPACE ALLOCATION BYPASSED DSN=dsname**
- Reason:** The current backup or input disk contain an extent of a multi-volume VSAM cluster which was flagged as a "candidate space". A candidate space is created when a multi-volume SMS cluster is allocated with a storage class with the "guaranteed space" attribute and real volume serials are specified. SMS preallocates extents on the volume, but they are not part of the active extents of the cluster until it expands and uses them.
- Action:** The candidate space will be bypassed and will not be allocated. FDR does not restore VSAM candidate volumes, either regular or candidate-space.
- FDR165 AIX RESTORE DEFERRED BASE CLUSTER NOT YET RESTORED AIX=name**
- Reason:** During an ABR restore of a multi-volume VSAM cluster with an alternate index (AIX), ABR detected that the AIX was on the current backup but the base cluster has not yet been restored, probably because it is on a different disk volume whose backup has not yet been processed. The AIX cannot be defined until the base cluster has been completely restored. ABR will bypass the AIX but will remember that it did so. After all other backups requested in this ABR restore have been processed, ABR will reprocess the backups containing the bypassed AIXs and attempt to restore them again (this may require remounting backup tapes). If the base cluster of the bypassed AIX was successfully restored from the other backups, the restore of the AIX will usually be successful. If the message appears a second time for the same AIX on the same backup, or the AIX restore fails for another reason, contact Innovation for assistance.
- FDR169 SMS NVR RECORD ERROR REASON=1**
- Reason:** On a SMS-managed volume, the NVR (Non-VSAM Record) in the VVDS had a length exceeding 256 bytes.
- Action:** Contact Innovation for assistance.
- FDR200 BLOCK DROPPED--synadaf--datablock**
- Reason:** A permanent BSAM read or write I/O error has occurred on a backup data set. "synadef" is descriptive error text provided by the IBM SYNADef macro. "datablock" is the first 20 bytes of backup data block read or written, in hex, if available. An IBM IOS000I I/O error message may also appear in the job's joblog. Tape I/O Errors in [Section 80](#) gives more details.
- DUMP:** If on reel-to-reel tape (3420-type tapes), FDR issues an EOVS to force a new output volume to be mounted. The blocks which were being written to the tape having the error will be rewritten as the first blocks on the new reel. If this tape switch occurs more than 20 times, FDR will terminate with a U0200 ABEND.
- On cartridges drives, because of the control unit buffer, FDR will issue the U0200 ABEND immediately after the first unrecoverable tape error. FDR will also terminate immediately if the

100.03 CONTINUED . . .

error occurs during the write of the control records at the beginning of the backup.

RESTORE: Any block in error is bypassed. If the error is due to I/O errors during the dump to a reel-to-reel tape, the block in error was rewritten on the next tape volume; if it is read successfully the data will be properly restored (Logical data set restore may fail on this condition). FDR will drop up to 20 blocks before terminating with a U0200 abend. If any track images were never successfully read from any tape volume, Message FDR366 (and possibly Message FDR155) will be issued to document this; if neither is issued then the restore was successful despite the I/O errors.

The block number in "synadef" will indicate the number of the tape block dropped. MAXERR= can increase or decrease the number of errors FDR will bypass. The keyword TAPERRCD=NO will suppress the condition code or ABEND that normally occurs at the end of the dump or restore.

Note: On reel-to-reel tapes (3420-type) MVS device swapping (the SWAP console command) should not be allowed for FDR backups. SWAP may be permissible for cartridge tapes. See Tape I/O Errors in [Section 80](#).

Note: Innovation offers two products which can dramatically reduce tape errors: FATS (FAST ANALYSIS OF TAPE SURFACES) will analyze tapes for defective areas and FATAR (FAST ANALYSIS OF TAPE AND RECOVERY) will display and correct data on data checked tapes.

FDR201 FORCE END OF VOLUME TAKEN--BLOCK WILL BE REWRITTEN

Reason: During dump processing a tape I/O error caused a tape volume switch to occur. See Message FDR200 for details. If BSAM caused the volume to switch during error recovery before FDR was able to force the EOv, this message will not be issued. If MVS device swapping (SWAP) was allowed to occur for cartridge drives, the message will not be issued (the only way to tell if a SWAP occurred is to observe in the joblog that the tape was dismounted from a different drive than that on which it was initially mounted).

FDR203 PREMATURE TAPE END OF FILE DSN=dsname

Reason: FDR restore detected an end-of-file on the backup data set without encountering the FDR trailer record. The dump may not have been successful. On a FDR full volume or data set restore, FDR will ABEND with a U0201; however, on a data set restore FDR will stop reading the backup when it has restored all tracks required for the data sets selected, so it may not issue the FDR203 or U0201 ABEND0.

The following are possible causes:

- The backup that created this file did not complete successfully.
- The backup run completed successfully, but the backup file was not correctly cataloged.
- The JCL for the restore specifies volume serial numbers for the backup file, but the last volume has been omitted, or the volume serials are out of order.

Action: In the last two cases, you can do the restore by specifying the correct volume serial numbers on the DD statement for the backup file.

FDR204 TAPE BLOCK LENGTH CHECK--BLOCK BYPASSED

Reason: 1. FDR stores the length of each block written to the backup in the first two bytes of the block. The length of the block read did not match the length stored in the first two bytes.
2. An error occurred during decompression of a block compressed by FDR software compression.

Action: An FDR mini-dump is printed displaying the beginning of the block. FDR will continue processing, bypassing this block. One or more tracks of data may have been lost. Message FDR366 will detail the tracks lost, if any. For a data set restore, unless message FDR155 is displayed, no tracks necessary for this restore were lost and the restore is successful. If Message FDR204 occurs many times, the cause is probably that you are restoring from a tape that was created by using a utility program (such as IEBGENER) to copy an FDR backup file.

Most utilities cannot copy FDR backup tapes. FDRTCOPY, FDRTSEL or FATAR must be used to copy FDR backup tapes. Other causes include transmission of an FDR backup over a network, or some channel extenders used for attaching tapes.

FDR205 TAPE BLOCK CROSS-CHECK ERROR

Reason: FDR encountered an error deblocking one of the tape blocks into track and disk block images.
Action: A mini dump is printed displaying the beginning of the bad block. A U0204 ABEND is issued if MAXERR= errors is exceeded.

FDR206 TAPE BLOCK COUNT ERROR

Reason: The number of blocks read from the tape did not match the block count in the tape trailer label.
Action: An FDR mini-dump is printed displaying the DCB and registers. Register zero (0) contains the block count from the tape trailer label, which reflects the number of blocks written to the tape volume during the backup. The fourth word of the DCB contains the count of blocks read

100.03 CONTINUED . . .

during the restore. FDR will ignore the error and continue processing. If tracks necessary to do the restore are lost, FDR will detail them in Message FDR366.

FDR207 TAPE CROSS-REFERENCE TABLE IS MISSING DSN=*dsname*

Reason: FDR could not find the track table on the backup data set "*dsname*". The backup may have been corrupted or overwritten.

Action: The restore will continue without the cross-reference check. Message FDR366 will not be issued even if data tracks were lost.

FDR210 I/O ERROR ON BACKUP DD=*ddname*

Reason: An I/O error occurred writing a FDR backup to the backup data set identified by "*ddname*". The message is followed by I/O diagnostic messages which identifies the error.

Action: Rerun the backup, using a different output tape and/or tape drive. Ask your tape hardware service representative to investigate the cause of the error.

100.03 CONTINUED . . .

FDR211 FDR ERROR ON DD=ddname REASON=reason

- Reason:** FDR InstantBackup discovered an error when attempting to initialize "ddname". "reason" can be:
- 8 FDR.USE.UNIT NOT FOUND** - the disk device specified by DSN=FDR.USE.UNITuuuu on the DISKx DD statement was not found in your I/O configuration.
 - 9 FDR.USE.UNIT IS ONLINE** - the disk device specified by DSN=FDR.USE.UNITuuuu on the DISKx DD statement is online. It must be offline to MVS.
 - A FDRINSTANT VOLSER MISMATCH** - the volume serial of the offline disk device being dumped does not match the volume serial of the online primary volume.
 - B INSTANT TRIAL HAS EXPIRED** - FDR InstantBackup trial has expired. Contact Innovation for assistance.
 - C FDR INSTANT UNIT NOT FOUND** - the disk device containing the FDR InstantBackup volume image was not found in your I/O configuration. For PGM=FDR or FDRDSF, if you specify BCV=, PPRC=, FCOPY= or SNAP=, this can also occur if you are trying to do Instant backups and non-Instant backups in the same step.
 - D FDR INSTANT PATH GROUP ERR** - FDR called the IBM module IECDVPTH to establish path groups for the offline disk, but the return code was non-zero. The backup will be completed but there may be some performance degradation.
 - E INSTANT BKUP NOT CATALOGED** - The latest backup generation and cycle as recorded in the ABR Model DSCB is not cataloged.
 - F INSTANT BKUP ALREADY DONE** - The latest backup generation and cycle as recorded in the ABR Model DSCB is currently cataloged to a tape or a disk other than the source disk itself, so the backup has already been moved. This message is normal if you restart the ABR backup step after some backups have been successfully moved to tape; it does not set any error codes and processing on the volume is bypassed.
 - G INSTANT BKUP UNIT MISSING** - A SNAP/SPLIT/PSPLIT/FCOPY was previously done, but FDR cannot locate the UCB for the device which was the target of that operation (which contains the volume image).
 - H INSTANT BKUP GEN/CYC ERROR** - The generation/cycle of the backup on the offline volume is not the same as current generation/cycle as recorded in the ABR Model DSCB.
 - J INSTANTBKUP TYPE= MISMATCH** - The TYPE= parameter used on the SNAP/SPLIT/PSPLIT/FCOPY operation was different from that specified on the subsequent DUMP statement.
 - K FDRINSTANT LABEL I/O ERROR** - An I/O error occurred reading the volume label from the offline volume.
 - N FDR INSTANT I/O ERROR** - An I/O error occurred while performing the InstantBackup.
 - P FDR INSTANT LABEL NOT FDR1** - when you do a SNAP/SPLIT/PSPLIT/FCOPY, FDR InstantBackup changes the VOL1 volume label on the output volume to FDR1 so it cannot be brought online by accident. But during the backup phase of FDR InstantBackup, the volume label was not FDR1 as expected.
 - Y UNABLE TO SPLIT - SPLIT RETRY COUNT EXCEEDED** - because of conditions within the Symmetrix, FDR was forced to re-establish and re-split the BCV a number of times, but the retry limit for this re-establish/re-split was reached before a valid split BCV was created.

Action: The backup is terminated (except for reason D). Many of these errors indicate that the FDR InstantBackup jobs have not been run in the correct order or with the correct parameters. Correct the error, if possible, and rerun. Otherwise, contact Innovation Technical Support. If you attempted to do a SPLIT/SNAP/PSPLIT/FCOPY, but it failed for some reason, then the DUMP step will get this message with reason E or F. Check the output of the SPLIT/SNAP/PSPLIT/FCOPY step to determine the cause of the error. However, reason F is normal if the DUMP step fails after processing some volumes, and you resubmit it; ABR will get reason F for the SNAP/SPLIT/PSPLIT/FCOPY volumes it has already processed and will bypass them automatically.

100.03 CONTINUED . . .

- FDR212** *program* **ERROR UNIT=uuuu REASON=reason**
Reason: FDR InstantBackup had an error attempting to process a volume on the device address shown as "uuuu", for the reason shown. Reason codes include:
- K** **EMC CODE nn - reason** - a BCV operation was requested, but the BCV request to the Symmetrix subsystem failed with the indicated reason code "nn". "reason" is a brief text explanation of the reason.
 - L** **CANNOT SPLIT VOLUME PAIR NOT ESTABLISHED** - a SPLIT operation was requested, but the primary volume was never ESTABLISHED with a FDR InstantBackup volume or the has been reused with a different primary volume.
 - U** **ERROR INVALIDATING VOL1 LABEL** - an I/O error occurred when FDR InstantBackup was attempting to modify the volume label of the BCV to insure that it cannot come online accidentally.
 - V** **CANNOT RESTORE VOLUME IS NOT SPLIT** - a request to restore from a captured volume image on a FDR InstantBackup volume failed because that volume is currently ESTABLISHED to a volume.
 - W** **BCV WAS SPLIT FROM ANOTHER STANDARD VOLUME** - a request to restore from a captured volume image on a BCV failed because that volume was most recently SPLIT from a different primary volume.
 - X** **ERROR RESTORING STANDARD VOL1 LABEL** - an I/O error occurred when FDR InstantBackup was rewriting the volume label on the target volume after restore from a FDR InstantBackup volume.
 - Y** **UNABLE TO SPLIT - SPLIT RETRY COUNT EXCEEDED** - the disk subsystem was unable to accept a SPLIT request. FDR retried the request but it continued to fail until the retry limit was exceeded.
- Action:** FDR InstantBackup will bypass this volume and continue processing.
- FDR218** *program* **UNIT=uuuu WAS stat UNIT=bbbb**
Reason: A BCV in an EMC Symmetrix or a ShadowImage pair in a HDS 7700E was used for FDR InstantBackup. This message indicates the status of the BCV volume as a result of the operation performed (RE-ESTABLISH, SPLIT or RESTORED). "uuuu" is the S/390 address of the primary volume and "bbbb" is the address of the FDR InstantBackup volume. "stat" is text which indicates the volume status:
- RE-CONNECTED TO** RE-ESTABLISH was executed after the dump due to BCV=(USE,RET) or PPRC=(USE,RET)
 - SPLIT FROM** SPLIT was successfully executed due to BCV=USE or PPRC=USE
 - RE-SPLIT FROM** SPLIT was executed due to BCV=USE or PPRC=USE but the target was already split from the primary, so FDR RE-ESTABLISHED it to refresh the target and SPLIT it again
 - RESTORED FROM** the online volume was restored from the FDR InstantBackup volume due to an ABR restorererequest
- FDR219** **FDR VOL=vvvvvv IS BEING DUMPED FROM OFFLINE UNIT=uuuu**
Reason: FDR InstantBackup was invoked. The online volume "vvvvvv" is being dumped from an offline copy on device address "uuuu". FDR InstantBackup is an extra-cost option to FDR.
- FDR224** **FDR VOL=vvvvvv IS BEING DUMPED FROM ONLINE UNIT BECAUSE SNAP/SPLIT/PSPLIT/FCOPY NEVER DONE**
Reason: A FDRABR DUMP is being executed with SNAP=USE, PPRC=USE, FCOPY=USE and/or BCV=USE but the identified volume is not eligible for FDR InstantBackup processing. It will be issued for every volume in the backup step which has not had a SNAP, PSPLIT, FCOPY or SPLIT executed against it successfully.
This message will be issued for all normal volumes (volumes against which you want regular backups, not FDR InstantBackups) that are processed in the same step with InstantBackup volumes, unless you have erroneously attempted to do a SNAP, PSPLIT, FCOPY or SPLIT of those ineligible volumes.
The message is also issued for all eligible volumes for which SNAP, PSPLIT, FCOPY or SPLIT has never been done, indicating that you want to backup those volumes normally.
- Action:** The volume will be backed up by normal ABR backup processing, but it will not be a point-in-time backup.
If the message identifies a volume which you did not intend to backup with FDR InstantBackup, it can be ignored.
If you **did** intend to use FDR InstantBackup on this volume, check your SNAP/SPLIT/PSPLIT/FCOPY step to be sure that it was included.
Note: if you attempted a SNAP/SPLIT/PSPLIT/FCOPY but it failed for some reason, ABR will fail that backup with message FDR211 REASON=E or F assuming that you wanted a point-in-time backup. Check the output of the SNAP/SPLIT/PSPLIT/FCOPY step to determine the failure.

100.03 CONTINUED . . .

- FDR225 FDRINST DID NOT *action* OFFLINE type UNIT=uuuu - OFFLINE BACKUP IS IN USE**
Reason: A FDRABR DUMP is being executed with SNAP=(USE,REL), PPRC=(USE,RET), FCOPY=(USE,REL) or BCV=(USE,RET) (or both) but at the end of the backup, ABR could not release the copied volume or RE-ESTABLISH the BCV because another task has the backup data set enqueued, probably because of a restore.
Action: You should RE-ESTABLISH the BCV or release the copied volume later when the backup data set is no longer in use.
- FDR230 *program* ERROR UNIT=uuuu REASON=reason**
Reason: FDR InstantBackup had an error attempting to special hardware facilities to get a point-in-time image of a volume for FDR InstantBackup use. "uuuu" is the device address of the target disk and "reason" indicates the cause. Note that some reason codes have several versions depending on the hardware facility which failed. Reason codes include:
- 1 **MOUNT STATEMENT NOT SPECIFIED** - There were no MOUNT statements following a SNAP/SPLIT/PSPLIT/FCOPY statement. MOUNT statements are used to define the source volume and target device for the FDR InstantBackup copy.
 - 2 **MOUNT STATEMENT NOT FOUND** - some control statements were provided after the SNAP/SPLIT/PSPLIT/FCOPY statement, but none of them were MOUNT statements. Only MOUNT statements can follow that statement.
 - 3 **xxxx UNIT= NOT SPECIFIED IN MOUNT STATEMENT** - a MOUNT statement does not contain the operand SNAPUNIT=, PPRCUNIT=, FLASHUNIT= or BCVUNIT=. It is required.
 - 4 **SOURCE AND type UNIT ADDRESS ARE EQUAL** - the unit on which the volume specified by VOL= is mounted is the same as the unit specified by xxxxUNIT=. A volume cannot be copied to itself.
 - 5 **INCOMPATIBLE SOURCE AND type UNIT DEVICE TYPES** - the device type of the unit on which the volume specified by VOL= and the unit specified by xxxxUNIT= are not the same. You cannot snap from a 3380 to a 3390 or vice versa.
 - 6 **type UNIT NOT FOUND** - the device address specified by xxxxUNIT= or was not found in the current I/O configuration. Be sure you specified a 4-digit hex address; add a leading zero if required.
 - 7 **type UNIT NOT DASD** - the device address specified by xxxxUNIT= was defined as a type other than disk in the current I/O configuration.
 - 8 **type UNIT IS NOT OFFLINE** - the disk address specified by xxxxUNIT= or was not varied offline to MVS. If you are sure that you want to use this device as a FDR InstantBackup target, vary it offline. Note that FDR InstantBackup will completely overlay that target volume and all existing data on the volume will be lost.
 - 9 **UNKNOWN DEVICE OR DEVICE NOT type CAPABLE** - either the source volume specified by VOL= or the target volume specified by xxxxUNIT= is capable of the special hardware feature requested (e.g, "ECAM" is the special command set used to perform Snapshot).
 - 9 **UNABLE TO OBTAIN SYMMETRIX DEVICE NUMBER** - FDR was unable to determine the MVS device address of the BCV assigned to the online volume being dumped. Possibly it is not in the I/O configuration or is in a different Symmetrix. Try specifying BCVUNIT= on the MOUNT statement.
 - 9 **UNABLE TO OBTAIN PPRC DUPLEX NUMBER** - FDR was unable to determine the MVS device address of the ShadowImage volume assigned to the online volume being dumped. Possibly it is not in the I/O configuration or is in a different disk subsystem.
 - A **UNEQUAL SOURCE/TARGET CONTROLLER FRAME SERIAL NO.** - the source volume specified by VOL= and the target volume specified by xxxxUNIT= are not both in the same disk subsystem (for FlashCopy, they must be in the same *logical* subsystem).
 - A **SYMMETRIX QUERY FAILED** - a query to the EMC Symmetrix to get BCV information failed.
 - B **SNAP STATUS INDICATES FAILURE** - Contact Innovation for assistance.
 - B **UNKNOWN DEVICE STATE FLAG=xx** - Contact Innovation for assistance.
 - C **SNAP STATUS INDICATES TIME OUT** - an internal timeout occurred in the hardware. If it is repeatable, contact Innovation for assistance.
 - D **SNAP STATUS INDICATES INVALID EXECUTE ATTEMPTED** - Contact Innovation for assistance.
 - E **SNAP STATUS INDICATES COUNT FIELD INCOMPATIBILITY** - Contact Innovation for assistance.
 - F **I/O ERROR DURING PROCESSING** - an unexpected I/O error occurred while processing the special hardware request. The error is documented by another message.
 - G **MESSAGE ID=mmmm STATUS CODE=ssss** - ECAM message number "mmmm" was issued to the Snapshot device. The message was accepted but it returned status code "ssss". Contact Innovation for assistance.
 - H **MESSAGE ID=mmmm COMPLETION CODE=cccc REASON=rrrr** - ECAM message number "mmmm" was issued to the Snapshot device. The message was rejected with completion code "cccc" and reason code "rrrr". Contact Innovation for assistance.

100.03 CONTINUED . . .

- I MESSAGE ID=mmmm PARMID=pppp COMP CODE=cccc REASON=rrrr** - ECAM message number "mmmm" was issued to the Snapshot device. The message was rejected because of a problem with message parameter "pppp" with completion code "cccc" and reason code "rrrr". Contact Innovation for assistance.
- J TARGET UNIT HAS LESS CYLINDERS THAN SOURCE DEVICE** - The source volume specified by VOL= has more cylinders than the target volume specified by SNAPUNIT=. The target volume must be equal to or larger than the source volume in size.
- K UNDEFINED SNAPSHOT STATUS** - - unexpected response from the Snapshot device.
- L UNDEFINED SNAPSHOT MESSAGE COUNT-** - unexpected response from the Snapshot device.
- M EXPECTED MSGID=nnnn . RECEIVED MSGID=nnnn-** - unexpected response from the Snapshot device; the wrong ECAM message was received.
- N INVALID FDRINSTANT VOL1 LABEL OR GEN/CYCLE MISMATCH** - FDR changed the ID of the volume label of the target disk to FDR1 so that it can't be accidentally be brought online. When the point-in-time image was dumped, the label was no longer FDR1, or the ABR generation and cycle created during the SNAP/SPLIT/PSPLIT/FCOPY was not the same as that currently recorded on the target. Contact Innovation if you do not want InstantBackup to modify the volume label.
- O PRIOR FDRABR INSTANT BACKUP PENDING DSN=dsname** - a SNAP/SPLIT/PSPLIT/FCOPY was not done because a previous SNAP/SPLIT/PSPLIT/FCOPY had been done, but you have not yet executed the DUMP to move the point-in-time backup to tape.
- P PRIOR FDRABR INSTANT BACKUP IS LOST DSN=dsname** - - a SNAP/SPLIT/PSPLIT/FCOPY detected that a previous SNAP/SPLIT/PSPLIT/FCOPY had been done, but you have not yet executed the DUMP to move the point-in-time backup to tape. Since you replied YES to the FDRW02 message or specified CONFMESS=NO, the uncompleted backup is overlaid and is no longer usable.
- Q TOO MANY WRITES PENDING** - - while FDRCOPY was attempting to use the EMCCOPY internal copy function on an EMC Symmetrix, the Symmetrix reported that an internal limit of pending writes (from cache to disk) was exceeded. The copy will continue using normal read/write I/O.
- Q CONTROLLER DOES NOT SUPPORT LRE/WFT** - Contact Innovation for assistance.
- R TARGET VOLSER DOES NOT MATCH SOURCE** - - the operand VERIFYVOLSER=YES was specified, but the current volser of the offline target does not match the volser of the online source volume.
- R DEVN/UNIT LOCATION ROUTINE ERROR FUNC=function** - Contact Innovation for assistance.
- S PPRCUNIT IS NOT PAIRED WITH PRIMARY** - the device specified by PPRCUNIT= is not currently in a PPRC session with the primary volume. You did not issue a previous (C)ESTPAIR command to establish the session.
- S ANTRQST IS NOT INSTALLED AT LEVEL 4-** you do not have the proper level of the IBM ANTRQST API to support FlashCopy installed. or contact Innovation for assistance.
- S ANTRQST REQUEST BUFFER EXHAUSTED** - an error occurred while invoking the IBM FlashCopy API. Contact Innovation for assistance.
- S SOURCE VOLUME NOT IN FLASHCOPY STATUS** - an error occurred while invoking the IBM FlashCopy API. Contact Innovation for assistance.
- S FLASHCOPY TARGET AND SOURCE ALREADY PAIRED** - the FlashCopy source and target volumes are each already paired with a different volume. FlashCopy cannot be used on this pair until those FlashCopy sessions are terminated.
- S FLASHCOPY UNEXPECTED QUERY DATA RETURNED** - an error occurred while invoking the IBM FlashCopy API. Contact IBM to be sure you have the latest maintenance for FlashCopy and contact Innovation for assistance.
- S FLASHCOPY ESTABLISH FUNCTION FAILED** - an error occurred while invoking the IBM FlashCopy API to establish a FlashCopy session. Contact Innovation for assistance.
- T FLASHCOPY TARGET AND SOURCE MIXED STATES** - either the FlashCopy source or target volume is already paired with a different volume. FlashCopy cannot be used on this pair until that FlashCopy session is terminated.
- S FLASHCOPY ESTABLISH FAILED FOR SOURCE** - an error occurred while invoking the IBM FlashCopy API to establish a FlashCopy session. Contact Innovation for assistance.
- T ANTRQST - PPRC REQUEST FAILURE** - an error occurred while invoking the IBM PPRC API. Contact Innovation for assistance.
- U SECURITY ACCESS DENIED TO DSN=FDR.USE.UNITxxxx** - FDR InstantBackup security checking was enabled, and the userid under which the FDR InstantBackup request was run does not have UPDATE authority to the data set FDR.USE.UNITxxxx where "xxxx" is the device address of the offline target volume specified by the xxxxUNIT= operand.

Action: In most cases, the volume will be bypassed. If additional MOUNT statements were specified, additional volumes will be processed.

100.03 CONTINUED . . .

- FDR231** *program* **COMPLETED SUCCESSFULLY - VOL=vvvvvv**
Reason: FDR InstantBackup has successfully used special hardware facilities to take an "instant" point-in-time image of the indicated volume.
- FDR232** *program action*
Reason: "action" may be:
RELEASED OFFLINE SNAPUNIT=uuuu - FDR InstantBackup was executed with SNAP=(USE,REL) to take a backup of the offline copy of an online volume created by FDR using Snapshot. Because of the REL parameter, FDR instructed the IBM RVA (RAMAC Virtual Array) or Storagetek SVA/Iceberg to release all of the tracks occupied by the offline device "uuuu". This reduces the NCL (Net Capacity Load) on the RVA/IcebergSVA/Iceberg by deleting the copied tracks which are no longer required.
RELEASED OFFLINE FLASHUNIT=uuuu - FDR InstantBackup was executed with FCOPY=(USE,REL) to take a backup of the offline copy of an online volume created by FDR using FlashCopy. Because of the REL parameter, FDR instructed the IBM 2105 ESS to withdraw the FlashCopy session to the offline device "uuuu"
DISBANDED PATHS FROM OFFLINE UNIT=uuuu - FDR InstantBackup will establish a channel path group to the office device for performance. When the dump was complete, the channel path group to the offline disk was disbanded.
- FDR270** **FDRCONVT OPTIONS IN EFFECT: options**
Reason: Issued to indicate the options in effect for program FDRCONVT.
- FDR271** **VOLUME operation RESULTS SUMMARY:**

VOLSER	FINAL STATUS	STORGRP	REASON FOR FAILURE
vvvvvv	status	storagegroup	reason

Reason: Issued by FDRCONVT to document the results of conversion for each volume processed. "operation" may be CONVERSION or SIMULATION. This message may be issued twice, once for volumes specified in the VOL= parameter and again for volumes dynamically allocated due to MULTIVOL=YES operand. "status" indicates the SMS status of the volume at completion of conversion. "reason" will indicate why the volume could not be converted, or blanks if conversion was successful.
- FDR273** **VOLUME vvvvvv ASSIGNED TO STORGRP=storagegroup**
Reason: Issued by FDRCONVT to indicate the SMS storage group assigned to a disk volume being processed. If no storage group is assigned, the "storagegroup" is 'NO STORG'. If the storage group was never determined, the "storagegroup" is '????????'.
- FDR274** **VOLUME operation status FOR VOLUME=vvvvvv**
Reason: the results of conversion for volume "vvvvvv". "operation" will be either CONVERSION or SIMULATION and "status" will be SUCCESSFULLY COMPLETED or FAILED
Action: If conversion or simulation failed, determine the reason for failure and rerun the job.
- FDR275** **VOLUME operation BEGINS FOR VOLUME=vvvvvv**
Reason: Issued by FDRCONVT to indicate that CONVERSION or SIMULATION has begun on volume "vvvvvv"..

100.03 CONTINUED . . .

- FDR276** **dsname status**
CATALOG=*catalog*
OLD CATALOG=*oldcatalog* **NEWCAT=***newcatalog*
VOLUMES=*vvvvvv vvvvvv ..*
STORCLAS=*storageclass* **MGMTCLAS=***managementclass*
DATACLAS=*dataclass*
DATASET SPANS MULTIPLE STORAGE GROUPS
Reason: The first form of FDR276 is issued by FDRCONVT to indicate the results on a specific data set. "dsname" may be
DSNAME=*dsname*
CLUSTER=*clustername*
COMPNT=*componentname*
AIX=*alternateindexname*
PATH=*pathname*
"status" indicates the final status of the data set:
CONVERTED TO SMS - conversion to SMS was successful.
CONVERTED TO NONSMS - conversion to non-SMS was successful.
ALREADY CONVERTED TO SMS - data set was already SMS-managed.
ALREADY CONVERTED TO NONSMS - data set was already non-SMS..
NOT ELIGIBLE - data set could not be converted. See message FDR277
The other forms of FDR276 may be issued after the first form, depending on the function performed. "catalog" is the catalog containing the data set; if FDRCONVT had to move the data set to a new catalog "oldcatalog" is the original catalog and "newcatalog" is the name of the catalog in which FDRCONVT cataloged the data set. "vvvvvv" is one or more volume serials containing the data set. "storageclass", "managementclass" and "dataclass" are the SMS classes assigned during a conversion to SMS; "(NULL)" indicates no class was assigned and "(NOT DETERMINED)" indicates that the class was not determined.
The last form of FDR276 is an informational message. A multi-volume data set was converted to SMS but the volumes on which it resides are not all on volumes in the same SMS storage group. SMS will never create this, but the data set will be usable.
- FDR277** **NOT CONVERTED REASON=***reason*
FUNCTION=*fn*, **RETURN CODE=***rc*, **REASON=***rsn*
Reason: The data set or cluster named in the preceding FDR276 message could not be converted by FDRCONVT for "reason". For errors in system calls (such as CATALOG), the second line of the message displays the function code "fn", the return code "rc" and/or the reason code "rsn". There are many possible "reason" values. Some common ones include:
DATA SET STORCLAS IS NULL - the ACS routines returned a null storage class for this data set, so it cannot be SMS-managed. You can specify the storage class via a SELECT with STORCLAS= but the ACS routines may still override it to null.
VVDS MANAGER ERROR - return/reason codes of 24 and 6 indicate that the VVR/NVR for the data set was not found in the VVDS.
ACS ROUTINE FAILURE - the reason code "rsn" probably corresponds to an IBM message IGCnnnnn; look it up in the IBM messages manual.
LOCATE FAILED FOR DATA SET - a catalog operation failed. The return and reason codes are documented under IBM message IGC3009I in the IBM messages manual.
DATA SET NOT CATALOGED - the data set to be converted was not cataloged and CATLG=YES was not specified.
Other reason values are possible.
Action: Processing continues with the next data set. Determine from the reason printed why the data set was not eligible. If necessary, call Innovation for assistance.
- FDR278** **VOLUME** *vvvvvv* **INELIGIBLE FOR CONVERSION: REASON=***reason*
Reason: The volume specified could not be converted to the desired status because it was not eligible.
Action: Determine from the reason printed why the volume was not eligible. If necessary, call Innovation for assistance.
- FDR279** **ddname DD MISSING**
Reason: A required DD statement was not found.
Action: Supply the missing DD statement.
- FDR280** **ddname OPEN ERROR**
Reason: A required data set failed to open.
Action: Determine the reason for the failure and rerun the job.

100.03 CONTINUED . . .

- FDR281 JOBCAT/STPCAT DD NOT ALLOWED**
Reason: A JOBCAT or STEPCAT DD statement was found in the JCL. JOBCAT and STEPCAT DD statements are not allowed with FDRCONVT.
Action: Use the INCAT parameter to specify catalogs to be searched for data sets catalogued outside the standard order of search.
- FDR282 SMS NOT ACTIVE ON SYSTEM**
Reason: The SMS address space is not active. FDRCONVT requires that SMS be active in order to run. This message may also occur if the operating system is not an MVS/XA or MVS/ESA system or if the Common VTOC Access Facility (CVAF) is not available.
Action: Start the SMS address space or run the job on a system where SMS is active.
- FDR283 NO VOLUMES TO PROCESS**
Reason: No volumes were online that met the specification in the VOL= parameter. FDRCONVT has nothing to do.
Action: Specify the correct volumes to be processed in the VOL= parameter.
- FDR284 VVDS CREATED ON VOLUME=vvvvvv**
Reason: FDRCONVT created a VVDS on a volume being converted to SMS status.
- FDR286 MASTER CATALOG NAME NOT FOUND**
Reason: The name of the master catalog could not be determined. FDRCONVT must determine the master catalog name in order to run.
Action: Correct any problems with the master catalog and rerun the job. If the problem cannot be determined, call Innovation for technical assistance.
- FDR287 ATTACH ERROR, RC=comp**
Reason: FDRCONVT could not successfully attach a required subtask. "comp" is the return code, in hex, from ATTACH.
Action: If the problem cannot be determined from the dump, call Innovation for technical assistance.
- FDR288 reason catalogname**
Reason: FDRCONVT encountered an error while trying to find a catalog specified in the INCAT parameter for the "reason" indicated. If the reason is LOCATE ERROR, the return code and reason code from CATALOG are displayed; these are documented under IBM message IDC3009I in the IBM Messages manuals.
Action: Correct the error. Specify the correct catalog if necessary. Contact Innovation for technical assistance if required.
- FDR289 ERROR OBTAINING type DEFINITIONS**
R15=nnnnnnnn, SFN=fn, RETN=rrrrrrr, RSN=ssssssss
Reason: FDRCONVT was unable to obtain the definitions for SMS classes; type is STORCLAS, MGMTCLAS or DATACLAS. R15 contains the return code from the SMS subsystem call. 'fn' is the SMS function code. 'rrrrrrr' and 'ssssssss' are the SMS return code and reason codes. The codes can be found in the IBM manual DFSMSdfp Diagnosis Reference.
Action: Correct the error. Contact Innovation for technical assistance if necessary.
- FDR290 type FACILITY CLASS NOT AUTHORIZED**
Reason: The authorization check for the specified RACF FACILITY class resource failed. "type" may be:
CONVERTVOL - the user was not authorized to resource STGADMIN.ADR.CONVERTV - required for FDRCONVT.
NCAT - the INCAT parameter was specified but the user was not authorized to resource STGADMIN.ADR.CONVERTV.INCAT.
BYPASSACS - the BYPASSACS parameter was specified but the user was not authorized to resource STGADMIN.ADR.RESTORE.BYPASSACS.
Action: Insure that the job has authority to the required RACF resources.
- FDR291 INCATNAME PARAMETER IGNORED -- CATLG=YES NOT SPECIFIED**
Reason: The INCAT parameter was specified without CATLG=YES when converting to SMS management. SMS-managed data sets must be catalogued in the catalog to which their high-level qualifier(s) are aliased.
Action: Recatalog the data sets to the correct catalog or specify CATLG=YES and rerun the job.

100.03 CONTINUED . . .

- FDR292 INVALID DSN/DSG NAME**
Reason: The DSN or DSG parameter did not begin with an alphameric or national character. The special form of the DSN or DSG parameter beginning with a period is not allowed with FDRCONVT.
Action: Revise the DSN or DSG parameter and rerun the job.
- FDR293 THRESHOLD PERCENT: LO=ll, HI=hh, PERCENT OF VOLUME IN USE=pp**
Reason: The high and low ALLOCATION/MIGRATION THRESHOLD values for the storage group to which this volume is assigned are displayed as well as the percentage of the disk volume that is in use.
- FDR294 FDRKWDPR INTERNAL ERROR**
Reason: An internal error was encountered while processing control statements.
Action: Contact Innovation for assistance.
- FDR295 INSUFFICIENT STORAGE**
Reason: Insufficient main storage was available for FDRCONVT to process all data sets on the volume.
Action: Increase the region size and rerun the job. If there are a large number of ICF VSAM components on the volume, increase the ICFCORE value.
- FDR296 WARNING: SPACE USED ON VOLUME EXCEEDS STORGRP THRESHOLD**
Reason: The percentage of the volume in use exceeds the upper value of the ALLOCATION/MIGRATION THRESHOLDS when converting to SMS management. The operating system may avoid allocating to this volume until the percentage of space used is decreased.
Action: Move data sets off of the volume until the percentage of the volume in use is less than the THRESHOLD value.
- FDR297 UNKNOWN PARM: parm**
Reason: An unrecognized parameter was specified on a control statement.
Action: Specify a valid parameter.
- FDR298 REQUIRED PARM parmname MISSING**
Reason: A required parameter was not found on a control statement.
Action: Supply the missing parameter.
- FDR299 VOLUME vvvvvv NOT FOUND**
Reason: Volume vvvvvv specified in the VOL= parameter was not mounted .
Action: Specify the correct volume and rerun the job.
- FDR301 program/function -- VER v.r/mmt -- INNOVATION DATA PROCESSING DATE=yyyy.ddd PAGE nnn**
Reason: This is the ABR page heading. See message FDR001 for details.

100.03 CONTINUED . . .

**FDR302 CONTROL STATEMENT ERROR NEAR REL LOCATION nn --
REASON x -- JOB TERMINATED**

Reason: An error was encountered during the processing of a user-supplied control statement. If "NEAR REL LOCATION nn" appears, the keyword or operand causing the error is at or near column "nn" on the input statement.

The error is defined by the reason code within the message. The failing statement is displayed immediately above.

Note: The expression "SELECT statement" in the reason codes below refers to SELECT and EXCLUDE statements. For FDRABR, it also refers to MOUNT, PROFILE and PROTECT statements.

The reason "x" can be:

- 1 A SELECT statement did not specify any operands. Control statement was blank after the Command name.
- 2 Command name on the first control statement was incorrectly specified. For example, FDRDSF must start with a DUMP, RESTORE, or PRINT statement. See the section in this manual for the program you are executing for details.
- 3 Operand on the first control statement was incorrectly specified. (Valid operands are TYPE=DSF, RESERVE, etc.).
- 4 Operand did not end with a blank or comma.
- 5 SYSIN data set was empty.
- 6 Expected continuation statement was not found. The previous statement ended with a comma and a blank.
- 7
 - a. On the first control statement, invalid or incompatible operands were specified. (Examples: RESERVE was specified but TYPE= was omitted, or FROM/TO was specified and TYPE= followed it.)
 - b. The TYPE= operand was omitted on a command that requires it (such as PRINT or COPY).
 - c. A control statement after the first was not a SELECT statement. A SELECT statement contained an invalid keyword. (Valid keywords are DSN=, VOL=, etc.).
- 8 An operand on a SELECT statement specified a blank or comma after the equal sign.
- 9 On the control statement printed above, one of the options exceeded its maximum length. A data set name exceeded 44 characters. A volume serial number exceeded 6 characters. The TAPEDD= option exceeded 1 character. The NOTIFY operand exceeded 7 characters. The DD= or NEWDD= operand exceeded 8 characters.
- B The DD= or NEWDD= operand specified a DDname which did not appear in the JCL.
- C Maximum number of SELECT statements was exceeded. The number of SELECT statements plus the number of commands in the PROTECT list was greater than the value specified in the MAXCARDS operand (default of 100 for ABR or 250 for FDRDSF).
- D ADATE= operand did not specify numeric data.
- E ADATE= operand is neither 5 digits (yyddd) or 7 digits (yyyyddd).
- F An operand did not specify numeric data (e.g., ADAYS=).
- G An operand did not end with a blank or comma or exceeded 15 digits (e.g., MAXCARDS=).
- H ADAYS= operand resulted in a negative year.
- I Keyword is invalid under the operation indicated. (e.g., GEN= was specified on a DUMP operation). Check manual for allowable keywords under this operation.
- J Control statement was completely blank. You can enter comment lines by placing an * (asterisk) in column 1.
- K A required operand was not specified on the preceding SELECT statement. On RESTORE TYPE=FDR, VOL= was not specified. On all other operations DSN, DSG, DD, ALLDSN or FROM/TO was not specified.
- L NEWG= began with a period.
- M COPY operand specified other than 1 or 2 on the SELECT Command, DSORG specified an invalid value, a date was improperly specified, the data set specified by CATDSN= is cataloged on more than 255 volumes or an invalid TAPEVOLGn= operand.
- O The ONLINE(ONLVOL) keyword is not supported on RESTORE TYPE=FDR operation. In a non-MVS system, ONLINE(ONLVOL) is not supported for any restore operation. The UNIT operand was specified on a non-MVS system.
- P ALLDSN was specified without ADATE on RESTORE TYPE=ARC. NEWN/G/I was specified with NEWNAME disabled. VOLG was specified on a RESTORE TYPE=ABR or TYPE=FDR.
- Q Keyword exceeded maximum value or was negative (e.g., GEN=10000 or CYCLE=64).
- R TYPE=xxx was specified multiple times.

100.03 CONTINUED . . .

- S** An operand on the SELECT statement was specified multiple times or was mutually exclusive with another operand. (e.g., DSN and DSG were both specified). DSN or DSG was specified on a MOUNT command.
- T** The operand BACKUP appeared on a SELECT statement read during a RESTORE TYPE=ARC operation or the operand ARCHIVE appeared on a SELECT statement read during a RESTORE TYPE=ABR operation. The operands BACKUP and ARCHIVE should only appear on restore requests written to the remote queue restore data sets (ABRREST and ABRARCH), respectively. Perhaps these two data sets were reversed.
- U** The TYPE=xxx operand was missing or invalid on the first statement. This operand is required for FDR, DSF, FDRCOPY, or ABR.
- V** An absolute track keyword was specified to FDRDSF (FROM/TO). The FROM/TO specification was in error. CYL and TRK did not specify numeric values or FROM value was higher than TO value, or absolute track operations have been disabled (NOABSTRK), or FROM/TO was specified on an EXCLUDE statement (DUMP only).

Action: Correct error and resubmit job. If you don't understand the error after reviewing the relevant sections of this manual, call Innovation for assistance.

Hint: if the control statements look good and you can't see any obvious reason for the error, check the JCL to be sure you are executing the right program for the statements you provided (PGM=FDR, PGM=FDRDSF, PGM=FDABR, etc.).

FDR303 CARD IMAGE - control statement image source

Reason: An input control statement is displayed by this message. "source" may be:
blank - from SYSIN
PARM ENTRY - from the JCL PARM= operand.
PROTECT LIST - from an ABR protect list
ALLOC LIST - from the ABR RESTORE ALLOCATION LIST.

FDR304 type function REQUEST FOR DDNAME=ddddddd,VOL=SER=vvvvvv,UNIT=uuuuuu

Reason: Describes the ABR operation that was performed.
 "type" is the type of operation and may be:
FDR--Full Volume Backup/Restore.
ABR--Incremental Volume Backup or data set restore from Volume Backup.
DSF--Manual Backup/Restore by data set.
ARC--Archive Backup, Application Backup or Restore data sets from either.
SCR--Superscratch.
 "function" will be **DUMP** or **RESTORE**. "ddddddd" is the DDname of the input disk (it may be DISKONLn if it was dynamically allocated by ABR). "vvvvvv" is the disk volume serial, and "uuuuuu" is the type and model of the disk as determined by ABR (e.g., 3390-3).

FDR305 direction TAPE DDNAME=dddddd,DSNAME=dsname, FILE=ffff,VOL=SER=vvvvvv,....

Reason: Describes the backup data set used in a ABR Dump/Restore operation. The DDname ("dddddd"), data set name ("dsname"), tape file sequence number ("ffff" - 0 for disk), and disk or tape volser are displayed. If over 5 volumes are involved, several lines will be printed. "direction" will be **TO** for backups and **FROM** for restores.

FDR306 function SUCCESSFULLY COMPLETED WITH DIAGNOSTIC MESSAGES

Reason: The ABR operation described by the message has been completed. If the text "WITH DIAGNOSTIC MESSAGES" appears, the dump/restore subtask detected non-terminating error conditions; you should check the messages printed on SYSPRINx. from the Dump/Restore subtask. "function" will be **DUMP**, **RESTORE**, **SIM** or **SCRATCH**.

FDR307 TAPE UNCATALOGED DSN=nnn...nnn

Reason: This ABR Volume Backup was a full-volume backup, so it started a new generation. Therefore, ABR uncataloged the backup tapes from all cycles of the n'th prior generation, where n is the number of generations to be retained for this disk volume, as recorded in the ABR model DSCB on the volume (ISPF panel A.I.8),

100.03 CONTINUED . . .

**FDR308 CATALOG FAILURE FOR BACKUP COND CODE=cc
MAXIMUM VOLSERS EXCEEDED**

Reason: A catalog request for an ABR backup data set has failed. "cc" is the catalog return code; common values are:

- 8** - ABR backup data set is already cataloged. Since the name of this backup duplicates a previous backup, this backup may not be usable. Some possible causes are:
 - The FDRABRM REMODEL command ([Section 50](#)) has been used incorrectly.
 - A full volume RESTORE or COPY under FDR or FDRABR has been done with VOLRESET=NO.
 - A Stand Alone Restore has been done.

These operations can back-level the generation and cycle number. The REMODEL command of program FDRABRM or the ABR ISPF dialogs can be used to update the model DSCB on the volume to the proper generation and cycle.

For an application backup (FDRAPPL), if the backup of the control file is a GDG, return code 8 can also occur if the GDG base has not been defined.

For an application backup (FDRAPPL), if the backup of the control file is a GDG, return code 8 can also occur if the GDG base has not been defined.
- 20** - the ABR catalog is full. Once you have expanded or reorganized the catalog you can manually catalog the backup using the information in the FDR305 message.

If the text "MAXIMUM VOLSERS EXCEEDED" appears a single ABR backup file has exceeded 19 tape volumes and cannot be cataloged. You may be able to reduce the number of volumes required by using higher-capacity tapes or using tape hardware compression (e.g., IDRC) or ABR software compression (COMPRESS=)

Action: If you cannot correct the error, contact Innovation for assistance.

**FDR309 DDNAME=ddddddd,VOL=SER=vvvvvv NOT ELIGIBLE FOR AUTO BACKUP I/O ERROR
COMP=cccc**

Reason: The disk volume named in the message has not been initialized for ABR processing; the ABR model DSCB was not found in the VTOC of the volume. If the text "I/O ERROR" is displayed, an error occurred attempting to read the ABR model DSCB and "cccc" is the return code from the OBTAIN SVC.

Action: To perform ABR processing (other than simulation and Application Backup) for this volume, you must initialize it for ABR processing with program FDRABRM ([Section 50](#)) or the ABR ISPF dialogs (panel A.1.8). If operand NOINIT is specified, ABR will allow a simulation of the ABR function (command SIM) to be done.

**FDR310 MAXIMUM NUMBER OF CYCLES/DATE EXCEEDED--
FDR DUMP FORCED ON VOL=vvvvvv**

Reason: An ABR Volume Backup requested an incremental backup, but a full-volume backup is being forced instead, because:

1. The number of cycles (incremental backups) within the current generation for this volume had reached the maximum (63).
2. The most recent full-volume backup of this volume had expired, according to the expiration date stored in the ABR Model DSCB. Specify DATEP=NONE on the DUMP statement to bypass this check and create the incremental backup.
3. The DUMP statement specified TYPE=AUTO, and the number of auto cycles created in the current generation was equal to or greater than the maximum specified in the ABR Model DSCB on the volume.
4. This was the first volume backup after a full volume restore with PGM=FDR, or with PGM=FDRABR using the TAPEDD option.
5. This was the first volume backup after the volume was initialized for ABR processing.

100.03 CONTINUED . . .

- FDR311** **FDR function DSN=dsname status
TO NEWN=newdsname
ON VOLSER=volser UNIT=device STORCLAS=sc
MGMTCLAS=mc DATACLAS=dc
CLUSTER=cluster [NEWC=newcluster**
Reason: The non-VSAM data set or VSAM component "dsname" was **SELECTED, DUMPED, RESTORED, COPIED, or MOVED**. This message will always be printed on data set operations, and will be printed on full-volume backups if PRINT=DSN is specified. Message FDR392 may be printed instead for data set dumps. On restore and copy/move operations "status" will also indicate:
- ALLOCATED** the data set was allocated by FDR; if absent, a preallocated dataset was overlaid.
 - CATALOGED** the data set was cataloged by FDR; if absent, either cataloging was not requested or the catalog request failed (see message FDR334).
 - LOGICAL** a logical restore was done, used when restoring to an unlike disk type, or when reblocking.
 - INSTANT** A hardware facility was used to quickly copy data in an FDRCOPY step. This occurs only if you are licensed for FDR InstantBackup. On an IBM RVA or StorageTek SVA/Iceberg with the Snapshot feature, Snapshot was used. On an EMC Symmetrix, an internal EMC copy function was used. Both the input and output volumes must be in the same RVA/SVA/Iceberg or EMC subsystem.
- If the data set or component is being restored to a new name, the line containing "newdsname" is displayed.
"volser" will define the volume serial of the disk to which it was restored, and "device" will identify the disk type and model (e.g., 3390-3). If the output data set is SMS-managed, the SMS classes assigned to the data set are displayed (management class and data class may be "(NULL)" if not assigned).
For a VSAM component, the line containing "cluster" will be displayed to identify the cluster to which the component belongs. If the cluster is being restored to a new name, "newcluster" is the cluster's new name.
- FDR312** **WARNING--FDR WAS SPECIFIED WITH DSNNAME CONTROL STATEMENTS--
STATEMENTS IGNORED**
Reason: SELECT statements were found following an ABR DUMP TYPE=FDR Command. They will be ignored and all data sets on the volume will be backed up. However, If EXCLUDE ALLDSN control statements are found, the volumes specified will be EXCLUDED from the full-volume dump if they are not referenced by DISKxxxx DD statements.
- FDR313** **fff TERMINATED BY OPEN EXIT ON VOL= vvvvvv**
Reason: The locally-written FDR volume open exit has terminated processing of the volume "vvvvvv". "fff" will be FDR, ABR, or CPK.

100.03 CONTINUED . . .

- FDR314** *program status1 dataset status2*
Reason: For an ABR Archive Backup or Superscratch ("program" is ABR), or an ABR Application Backup (APP) or an FDRCOPY MOVE (FDR), this documents the action taken on one non-VSAM data set, VSAM component, or VSAM cluster. It can also be produced on any type of restore if RECAT or VRECAT is specified.
 "dataset" will be:
 DSNAME=dsn for a non-VSAM data set
 CLUSNM=cluster for a VSAM cluster or
 VSAM COMP=component for a component of a VSAM cluster
 "status1" can be:
 ARCHIVED AND SCRATCHED - Archive Backup backed up and scratched the data set.
 ARCHIVED NO SCRATCH - Archive Backup backed up the data set but it was not scratched because SCRATCH=NO was specified.
 ARCHIVE DEFER SCRATCH - Archive Backup backed up components of a multi-volume cluster from one volume but deletion of the cluster is deferred until all volumes containing the cluster have been processed (they must all be processed in the same ABR step).
 WITH NO BACKUP - Superscratch deleted the data set.
 APPLICATION BACKUP OF - Application Backup backed up this data set.
 SCRATCHED - an FDRCOPY MOVE scratched the input data set after copying it to the output. Can also occur when VRECAT causes an existing VSAM cluster to be deleted.
 UNCATALOGED - an FDRCOPY MOVE or restore with RECAT caused an existing SMS-managed data set to be uncataloged
 "status2" can be blank or one of these:
 AND UNCATALOGED - Archive Backup, Superscratch or MOVE also uncataloged the scratched data set.
 AND UNCATALOGED RECAT FAILED - a VSAM cluster was archived with RECALL=YES and the attempt to catalog it as a non-VSAM data set for auto-recall failed, leaving it uncataloged. It cannot be auto-recalled.
 RECAT FOR RECALL - a data set was archived with RECALL=YES. The catalog entry has been updated to indicate that the data set is auto-recallable.
 RECAT FOR RECALL FAILED - a non-VSAM data set was archived with RECALL=YES but an error occurred attempting to recatalog the data set for auto-recall. The data set is still cataloged but is not auto-recallable.
Action: If the status includes RECALL FAILED or RECAT FAILED, the indicated data set is not auto-recallable. If you can delete any existing catalog entry for the data set, you may be able to use the RECATALOG function of FDRARCH ([Section 51](#)) to build the auto-recall catalog entry.
- FDR315** **DDNAME=ddn, VOL=SER=vvvvvv NOT ELIGIBLE FOR function**
MAXIMUM RUNS EXCEEDED
Reason: 1) The ABR Model DSCB on this volume does not indicate that the volume is enabled for ARCHIVE or **SCRATCH** (Superscratch) or the ABR Model DSCB does not exist.
 2) If MAXIMUM RUNS EXCEEDED is printed, more than 255 Archive Backups of this volume were done today; this limit is only enforced when the backups must be cataloged in the ABR catalog ([see Section 51](#)).
Action: This volume will be bypassed. If you intend to do Archive Backup or Superscratch on this volume, create or update the ABR Model DSCB with FDRABRM ([Section 50](#)) or ISPF panel A.I.8 ([Section 90](#)).
 If the operand NOINIT is specified, ABR will allow a simulation (SIM TYPE=) of the volume to continue.

100.03 CONTINUED . . .

- FDR316** **function status REQUESTED DSN/DSG=dsname VOL=vvvvvv**
Reason: A SELECT or EXCLUDE statement was input to an FDR program, but none of the data sets processed by that program matched the data set name, dsname mask, or data set group name ("dsname") on that statement. This is probably an error, since there must have been a reason for providing the statement. "vvvvvv" is the volser if specified by VOL= on the SELECT or if selected from the catalog by CATDSN=.
 If "status" is **DID NOT FIND**
 1) for backups, no data set on the input disks matched the name specification.
 2) for DSF restores, it means that no such data set was found on the input backup files.
 3) for ABR restores, it means that Volume Backup information or Archive Backup information could not be located. This can also be caused by a EXCLUDE preceding the SELECT for the same data sets. Another error (such as Message FDR325) may also cause the data sets not to be selected.
 If "status" is **EXCLUDED**, the message simply documents that the named data sets were excluded.
 If "dsname" is **MOUNT**, then an ABR MOUNT statement did not select any volumes for ABR processing.
 If an ABR full-volume dump was forced (see message FDR310), then SELECT statements were not processed, and Message FDR316 appears for every SELECT. However, all data sets on the volumes were dumped.
Action: If this is an error, correct and resubmit job. If you did not expect to get this message, there is probably some error in your SELECT statements; check them. One common error: if you have more than one SELECT statement, and an earlier one selects all the data set that would be selected by a later one, the later one will get the FDR316. For example,
 SELECT DSN=A.B.**
 SELECT DSN=A.B.C
 In general, put the SELECTs with the most restrictive selection first.
 This message will usually cause the FDR step to end with an ABEND or non-zero return code. If you expect to get this error (perhaps because some data sets may not always exist) and want to ignore it, specify SELTERR=NO on the DUMP statement.
Note: if the FDR316 message appears in the job log of a batch job and is followed by a system ABEND S213-04 **this does not indicate an error in ABR**. This is for a recall which was invoked by the ABR Data Set Not Found (DSNF) exit, which is entered whenever a program tries to OPEN a data set but the DSCB for that data set was not found in the VTOC of the allocate volume. The DSNF exit invokes ABR to recall the data set in case it was archived; however, if it was never archived or has expired the FDR316 message is printed to document that ABR tried to recall the data set but it was not recorded in the Archive Control File. This probably means that the data set name was misspelled by the user or the wrong volume serial was given. It can also be caused by trying to OPEN a VSAM cluster with a non-VSAM OPEN (looks for the cluster name instead of the component names in the VTOC).
- FDR317** **I/O ERROR ON DISK TRACK X'cccchhhh'**
I/O ERROR ON DISK - DDNAME ddname - SEEK ADDRESS ccccchhh
Reason: ABR encountered an I/O error while reading or writing the VTOC. The FDR mini-dump will be printed.
Action: ABR will continue processing the VTOC bypassing this track. Contact Innovation for assistance.
- FDR318** **DUMP COMPLETED WITH NO DATA SETS SELECTED**
Reason: A data set backup was made with no data sets being selected by either control statements or automatic functions.
Action: If this is an error, correct and resubmit.
- FDR319** **FDR OPERATION ABNORMALLY TERMINATED VOL=vvvvvv COMP CODE=Ssss Uuuuu**
Reason: An internal FDR subtask failed on volume "vvvvvv".
 If "sss" is non-zero (in hex), the task failed with a System Ssss ABEND. Consult IBM documentation for the meaning of the ABEND. Sx13/Sx14 ABENDs are OPEN/CLOSE errors and are accompanied an explanatory message in the joblog of the FDR job.
 If "uuuu" is non-zero (in decimal), the task failed with a User Uuuuu ABEND. The ABEND is probably generated by FDR; see [Section 100.08](#) for FDR User ABENDs. Most FDR User ABENDs are preceded by an FDR message.
Action: If the problem cannot be determined from the ABEND code, call Innovation for technical assistance. If the ABEND produced a storage dump, have it available.
 Only this volume is affected, FDR will continue to process other volumes. However, the step will end with an ABEND or non-zero return code to call attention to the error.
 If this is an ABR backup to tape, and there are more disk volumes to be processed, ABR will call for a fresh scratch tape on the output drive for the TAPEX and TAPEXX DD statements in use by the dump subtask that ABENDED.

100.03 CONTINUED . . .

- FDR320 UNABLE TO function REASON=x--DSNAME=dsname**
Reason: FDR was unable to **RESTORE, COPY or MOVE** the data set specified in the message. The reason code "x" may be:
- 1** The FDR restore subtask restoring this data set from a backup failed; see message FDR319, other FDR messages may also appear..
 - 3** I/O error reading a Format 2 or Format 3 DSCB from the output disk.
 - 6** Data set enqueue failed for this data set. DSNENQ=NONE will override this check.
 - B** Data set is unmovable.
 - C** Data set not found on the backup file. The user may have misidentified the backup (backup does not contain the data set), or it may be due to:
 - 1. The same dsname was specified on more than one control statement or the dsname is a component of previously requested ICF VSAM cluster.
 - 2. For ICF VSAM clusters, can occur if there was more than one VVR for a component of the cluster in the VVDS of the original disk when dumped (this is an error).
 - D** Data set has an F3 DSCB, but it could not be found on the backup.
 - E** Data set was being restored to a new name and the FDR security checking option (ALLCALL) is enabled but: 1) the user does not have READ authority to the original data set.
 2) the original data set is not protected by your security system (this is considered an error since the data set may have been backed up on another system where it was protected).
 - G** Unable to locate data set's F3 DSCB on target volume or F2 DSCB found for a non-ISAM data set.
 - J** Data set was RACF or password protected when it was backed up, but user has pre-allocated it without protection, or vice versa.
 - K** Data set has not reached its expiration date, and the operator replied "NO" to the FDRW03 message.
 - O** PRESTAGE was specified. Data set restore bypassed since data set was found.
 - P** Characteristics of the input data set and the pre-allocated output data set do not match. Examples of this:
 VSAM to non-VSAM PDSE to non-PDSE
 HFS to non-HFS Extended Format SMS to non-EF
 and vice-versa in each case; other combinations also cause errors.
 - S** A data set which must be SMS-managed (such as Extended Format (EF)) is being restored to a non-SMS volume.
 - T** FDRCOPY COPY/MOVE detected a copy/move of a data set on top of itself (same volser, same data set name).
 - U** Error updating ISAM Format 2 DSCB, or ICF VSAM component in the VTOC but not in the VVDS.
 - V** ICF VSAM cluster is being restored to a VSAM cluster which has different options or characteristics. Message FDR152 or FDR159 details the error.
 - W** Catalog inaccessible. Possible causes are:
 - 1. When restoring an ICF catalog, the JCL contained a STEPCAT or JOBCAT DD statement. This is not supported.
 - 2. When restoring an ICF catalog, the catalog being restored was defined in the master catalog, but it was not usable (probably did not actually exist on disk). You must remove the entry for the old catalog, probably by IDCAMS EXPORT DISCONNECT.
 - X** Data set security exit rejected this data set.
 - Y** NEWINDEX= failure or duplicate data set name generated. It may be that the new data set name does not meet IBM standards, or that the resulting name is greater than 44 characters.
 - Z** One or more data sets failed to restore using ALLDSN or DSN=mask. Check the subtask listing for errors.
- Action:** Take appropriate action and/or contact Innovation for assistance.

100.03 CONTINUED . . .

FDR321 UNABLE TO RESTORE REASON=x--DSNAME=dsname**Reason:**

The data set specified in the message could not be restored. The reason code "x" may be:

- 2** The ABR JCL does not include any DISKxxxx DD statement pointing to a disk volume of the same type (e.g., 3390) as the original volume of the data set.
- 3** The OPERATOR operand was specified, but the operator replied 'NO' or 'BYPASS' to the request for the backup tapes needed to restore the data set.
- 5** The disk volume which FDR was attempting to restore the data set could not be opened, or could not be dynamically allocated.
- 7** During a FDRCLONE restore, ABR extracted a volume serial and generation/cycle numbers stored in a cloned data set's catalog entry by FDRCLONE, but that backup was not cataloged in the ABR catalog.
- A** For an ABR restore from Volume Backups, the data set was found on disk and ABR is trying to get backup information from its DSCB, but no current backup information was recorded. Either the data set has never been backed up by ABR or it has been deleted and reallocated since its last backup. This can also occur if some other software has corrupted ABR's backup information in the DSCB. If you can identify an ABR backup which contains the data set, you can specify GEN=, CYCLE= and VOL= to restore it.
- B** For an ABR restore from Volume Backups:
 - 1) the data set is not cataloged and VOL= was not specified
 - 2) the data set is not on the volume identified by the system catalog (or by the VOL= operand)
 - 3) It was also not found in ABR's record of scratched and renamed data sets (the ABR SCRATCH catalog). The most likely causes are:
 - a) The DSN= or VOL= operand is incorrect.
 - b) The data set was archived.
 - c) ABR is not recording scratched and renamed data sets (ABR DADSM exit not installed)

If you can identify an ABR backup which contains the data set, you can specify GEN=, CYCLE= and VOL= to restore it.
- D** For an ABR restore from Volume Backups, ABR identified a specific backup containing the requested data set, but that backup was not cataloged in the ABR catalog. The backup may have expired or was part of a older generation that ABR has uncataloged. If COPY=n was specified, it may be that the specified copy was not created. If the backup file still exists but has been uncataloged, TAPEDD= can be used to restore the data set.
- E** For an ABR restore from Archive Backups, the backup file containing the data set has reached its expiration date (as recorded in the Archive Control File). If the backup file still exists, you can restore the data set by specifying TAPE=EXP on the RESTORE control statement.
- F** For an ABR restore from Volume Backups, the volume from which the data set was backed up, as specified by the VOL= operand or the catalog, is not currently initialized for ABR processing; or an I/O error occurred when reading the ABR model DSCB controlling the disk volume. If you can identify an ABR backup which contains the data set, you can specify GEN=, CYCLE= and VOL= to restore it.
- G** SELECT statement for a data set restore did not specify VOL=. This data set was not found in the system catalog, or in ABR's record of scratched data sets. See reason B for suggestions.
- H** The volume from which the data set was backed up, as specified by the VOL= operand or the system catalog, was not mounted so ABR could not locate the backup of the data set. If possible, mount the volume; If you can identify an ABR backup which contains the data set, you can specify GEN=, CYCLE= and VOL= to restore it.
- J** GEN= was specified without CYCLE=.
- K** TAPEDD=x was specified but no TAPEx DD was found, or TAPEDD was specified with TYPE=ARC.
- L** The data set was not found on the volume specified by the VOL= operand (if coded) or by the catalog or it was found but had no current ABR backup recorded in its DSCB. The data set was found in the ABR SCRATCH catalog, but with a different volume serial. This probably occurred because the data set was scratched and reallocated on a different volume. You can restore the data set by specifying VOL= the volume in the ABR SCRATCH catalog, or by specifying VOL= and CYCLE= (and possibly GEN=) for any volume and cycle on which the data set was backed up, or by specifying TAPEDD=.
- M** COPY=2 was requested for this data set, but the ARCHIVE control data set indicates that a second copy was not created during archival .
- N** The backup file containing this data set failed to OPEN, or dynamic allocation error. Check the job log for IBM error messages. The backup data set may have been written over.
- Q** RACF authorization error.

100.03 CONTINUED . . .

- R The backup file containing this data set was not an FDR-formatted backup; or an I/O error occurred reading the FDR Header Record.
- T A TYPE=ABR restore had a SELECT statement that specified DSG= or ALLDSN without either TAPEDD= or VOL=, GEN= and CYCLE=.
- V For a restore from Archive Backup or Application Backup, the backup file exceeded 5 tape volumes, but ABR was not able to find an entry for it in the ABR catalog. The catalog will contain the complete volume list, while the Control File only records the first 5).
- X A DSF absolute track restore was requested (SELECT FROM/TO) and cylinder or track numbers were specified which are invalid for the type of disk which was backed up.
- Y OLDBACKUP was specified. The data set does not contain as many levels of OLDBACKUP information as requested by the OLDB=nn operand, or the specified backup is no longer in the ABR catalog.
- Z An ABR data set restore was requested and a TAPEX DD statement was provided, but the backup containing this data set occupies more than 5 tape volumes. To restore this data set remove TAPEX DD and specify the DYN TAPE option. This code will also occur if you use the TAPED=x operand on an archive restore (TYPE=ARC) or application restore (TYPE=APPL).

Action: Take appropriate action for the problem described by the message.

FDR323 DDNAME=DISKdddd DOES NOT SPECIFY A DISK

Reason: DD statements whose DDnames start with DISK must point to a disk device.

Action: Correct and resubmit job.

FDR324 DDNAME=dddddddd IS MISSING OR IN ERROR REASON=x

Reason:

A failure has occurred for DDname specified in the message. In the reason text below "Control File" refers to the Archive Control File for Archive Backups, and to the Application Control File for Application Backups. The reason code "x" may be:

- 1 SYSIN DD statement missing.
- 2 I/O error on SYSIN data set or remote queue data set.
- 3 ARCHIVE DD statement missing and DYNARC was not specified.
- 4 Control File block size is zero. The file may have been allocated but never initialized. If so, use FDRARCH to initialize it.
- 5 I/O error on Control File.
- 6 The date of the last Archive or Application backup, as recorded in the Control File, is higher than the current system date; this might occur if you did backups while you had the system date set ahead for YEAR2000 testing. To override the ABR date check, specify PROT=NONE on the DUMP statement of next Archive or Application job using that Control File only. The Control File will be out of date sequence.
- 7 A block of the Control File did not contain ABR identifiers. The Control File may have been overwritten by another program. You may be able to restore a backup of the most recent Control File. Please call Innovation for assistance..
- 8 The Control File is full. You can reorganize or expand it with FDRARCH. If you are using the option of Application Backup to create the Control File, increase the space allocated.
- 9 The number of TAPEX DD statements exceeded 9.
- A No TAPEX DD statements were specifying tape or disk backup were found. ABR ignores DD DUMMY statements
- B The number of DISKxxxx DD statements plus any volumes selected by ONLINE, ONLVOL, and MOUNT was greater than the value specified in the MAXDD operand (default of 256).
- D A SYSPRINx DD statement was missing. ABR requires a SYSPRINx DD statement for each TAPEX DD statement.
- E A DISKx DDname was used more than once in the JCL for this step.
- F A FDR or FDRDSF backup specified DISP=MOD on the TAPEX DD card. Also occurs if the TAPEX DD refers to a subsystem data set (such as ExHPDM) via the SUBSYS= JCL parameter, and the RTC=YES or DCT=YES operand was also specified on the DUMP statement; these are not compatible. The disk volume associated with this tape will not be processed.
- I OPEN failed for the backup file. Check the joblog for IBM error messages.
- K LASTAPE was specified but the TAPEX DD specified more than one tape volume serial.
- L DYN TAPE was specified to dynamically allocate the backup medium but a TAPE# DD statement is in the JCL. Remove it; it is required for DYN TAPE dynamic allocation.

100.03 CONTINUED . . .

- M** COPY or MOVE was specified. The TAPEX DD statement for the target volume was not a disk or specified DD DUMMY.
- N** The TAPEX DD statement specified the same volume serial as the DISKx DD statement.
- O** Internal error detected by module FDRDSKTP. Contact Innovation for assistance.
- Q** An ABR dump specified DISK as the backup medium (TAPEX). ABR found that the volume to be dumped (DISKx) was the same as the backup medium (TAPEX). You cannot backup a volume to itself; it will be bypassed.
- R** ABR space allocation was not successful for a backup data set on any of the disk volumes specified by the TAPEX DD statement. Possible reasons are:
 - 1) The output disk volumes may be full or have insufficient space for this backup
 - 2) The backup data set name being allocated already exists on the output volume (possibly due to an aborted ABR run). Scratch the duplicate backup data set and run the backup again.
 - 3) the output disk volumes are SMS-managed (not supported for ABR output disks).
- T** User specified multiple TAPEX DD statements using the same last character (x). The first one specified DUMMY.
- U** Dynamic allocation or OPEN of the backup data set on a disk volume failed. Check the joblog for IBM error messages or the SYSPRINx printouts for FDR error messages.
- Action:** Take appropriate action for the problem described by the message.

FDR325 FDRABR NO DISK DD STATEMENTS FOUND

- Reason:** The ABR step did not identify any disk volumes to be processed. Disk volumes can be identified to ABR by:
- 1) DD statements having DDnames starting with DISK.
 - 2) The ONLINE operand (all online volumes)
 - 3) MOUNT statements
 - 4) The ONLVOL operand (requires that volumes be identified by SELECT statements with either VOL=, VOLG=, or CATDSN= operands.
- (3) and (4) can cause message FDR325 if none of the identified volumes are online. If this is a job which processes an ABR remote queue data set, the message can also occur if the remote queue is empty.
- Action:** Change the jobstream to properly identify the disk volumes to be processed.

FDR326 function ENDED WITH NO dsvol process

- Reason:** DSF, FDRCOPY, or ABR was not able to successfully process any of the data sets or volumes specified; or, no SELECT Commands followed a RESTORE, COPY or MOVE Command.
- "function" will be **RESTORE, COPY, MOVE** or **REORG**.
- "dsvol" will be **VOLUMES** or **DATA SETS**.
- "process" will be **RESTORED, COPIED, MOVED SELECTED**, or **REORGED**.
- Action:** Correct any specification errors and resubmit. If you wanted to select a VSAM cluster, make sure you specified the VSAM cluster name and not the component names.

FDR327 CATALOG FAILURE TO LOCATE TAPE DSNAME=dsname

- Reason:** A restore from ABR Volume Backups was requested (RESTORE TYPE=ABR). The required backup file was not found in the ABR catalog. If restoring an older backup with the OLDB=nn operand, it may be that the backup generation and cycle selected has been uncataloged by ABR, or expired by your tape management system.
- Action:** If the required backup tape is still available (has not been overwritten), even though it is not cataloged, restores can still be done from it using the TAPEDD=n operand, and providing a TAPEn DD statement specifying the ABR backup data set name, volume serials, and file number (LABEL=n). This information can be found in the FDR305 messages printed during the backup, or from other ABR reports.

100.03 CONTINUED . . .

- FDR328 UNABLE TO RESTORE REASON=reason--VOL=vvvvvv NVOL=vvvvvv**
Reason: ABR was unable to perform a full volume restore (RESTORE TYPE=FDR). "vvvvvv" is the original volume serial of the input volume and "nnnnnn" is the current volume serial of the output volume. "reason" may be:
1 - RESTORE VOL OR DISK(X) MISSING - ABR did not find a DISKxxxx DD statement pointing to the disk volume to be restored to.
2 - NO ABR MODEL OR NO BACKUPS - GEN= was omitted, but either the original volume was not online or an error occurred reading the ABR model from the volume AND no backups were cataloged for this volume in the ABR catalog. So, ABR cannot find backups for this volume. At a disaster site, be sure that the ABR catalog has been restored to its latest version.
3 - TAPEDD NOT MATCHING A TAPEX DD - the TAPEDD=x operand was specified. ABR could not find a TAPEX DD statement with the same character as specified in this operand.
4 - GEN SPECIFIED WITHOUT CYCLE
5 - GEN=CUR - NO BACKUPS CATALOGED GEN=CURRENT was specified but either no backup entries for this volume found on the ABR catalog or an I/O error occurred reading the catalog.
6 - SMS TO NON-SMS OR VICE-VERSA - A backup of an SMS volume is being restored to a non-SMS volume, or a backup of a non-SMS volume is being restored to an SMS volume. If "OVERRIDDEN" is displayed, then an SMS backup is being restored to a non-SMS volume using SMSPROT=NONE.
7 - FDRDRP DISASTER TRIAL HAS EXPIRED - your trial of FDRCLONE, which includes FDRDRP, has expired. Contact Innovation if you want to extend your trial or license FDRCLONE.
8 - INSUFFICIENT STORAGE SPECIFIED - a GETMAIN for a required work area failed. On ESA and OS/390 systems, the work area is above the 16M line. Increase the size of your above the line private area with REGION=nnM on the EXEC JCL statement.
- FDR329 DD=ddname1 IS A DUPLICATE VOLUME OF DD=ddname2**
Reason: The disk volume referenced by DD statement "ddname1" is the same as that referenced by "ddname2" ABR will process the volume only once.
- FDR330 LAST TAPE OR POOL OPTION FAILED ON TAPE DD=ddname REASON=reason**
Reason: ABR JCL options were used to request for the LASTAPE, POOLTAPE or POOLDISK features of ABR, but the request failed. "reason" can be:
1 - LASTAPE FILE FAILED TO OPEN - The LASTAPE file identified in the JCL was cataloged as having been created in a previous ABR run, but the attempt to OPEN it failed. It may have been overwritten. Check for IBM OPEN error messages in the joblog. ABR uncataloged the LASTAPE file so if you resubmit the job, ABR will ask for a scratch tape.
6 - POOLDISK DOES NOT SPECIFY DISK - must point to disk volumes.
9 - LASTAPE/POOLTAPE NOT A TAPE - must point to a tape device.
A - LASTAPE UNCAT FAILED COMP=nnnn - The request to uncatalog the LASTAPE data set failed with catalog return code "nnnn".
C - POOLDISK SORTING ERROR SWAREQ - while sorting the POOLDISK volumes by space available, an error occurred during a SWAREQ macro to access the volume list.
D - POOLDISK SORTING ERROR DEVTYPE - while sorting the POOLDISK volumes by space available, an error occurred doing a DEVTYPE macro on one of the volumes.
E - POOLDISK SORTING ERROR UCB - while sorting the POOLDISK volumes by space available, an error occurred accessing a UCB.
F - POOLDISK SORTING ERROR LSPACE - while sorting the POOLDISK volumes by space available, an error occurred doing an LSPACE macro to get the free space on the volume.
G - POOLDISK DOES NOT SUPPORT SMS - POOLDISK volumes must be non-SMS managed..
Action: You may be able to avoid reasons C through F by specifying POOLSORT=NO on the ABR DUMP statement. If cause cannot be determined, contact Innovation for assistance.

100.03 CONTINUED . . .

**FDR331 ARCHIVE BACKUP TERMINATED WITH HARDWARE ERRORS
[BAD FORMAT 4]**

Reason: ABR encountered disk or tape I/O errors during the physical dump of the data sets to be archived. Another message is printed in the SYSPRINx report to document the exact error. Since ABR would normally scratch these data sets after dumping them, a clean backup is imperative. Therefore, the archive run for this volume has been aborted. If BAD FORMAT 4 is printed, the cause was an invalid or duplicate Format 4 DSCB detected in the VTOC when preparing to scratch the data sets.

Action: Correct error and rerun job. If the error cannot be easily corrected, EXCLUDE the data set causing the error.

FDR332 function SCRATCH status ON DSN=dsname

Reason: An attempt to delete a non-VSAM data set **FAILED** or was **BYPASSED**. FDR deletes data set during an FDRCOPY MOVE, an Archive Backup (DUMP TYPE=ARC), or ABR Superscratch (DUMP TYPE=SCR). "status" of **FAILED** indicates that a SCRATCH SVC (for non-SMS data sets) or DELETE NVR or NONVSAM (for SMS-managed data sets) failed to scratch this data set. An FDR mini-dump is printed showing the registers and SCRATCH or DELETE parameter list. For a non-SMS data set, the SCRATCH status code appears in the mini-dump, in the first line of the storage display after the registers, in the second byte of the last word (right most word). Code 1 indicates that the data set was not found in the VTOC. Code 7 indicates that SCRATCH failed because the data set was in use; this can be avoided by using the DSNENQ= option. Code 2 or 8 indicates a security violation. For an SMS data set, Register 15 (R15) contains the DELETE return code, and the reason code is in the last 2 bytes of the 6th word of the first line of the storage dump after the registers, in hex. These codes can be found under Message IDC3009I in the IBM Messages manual after converting the codes to decimal. For Archive Backup, ABR has backed up the data set and recorded it in the Archive Control File. This data set may have been uncataloged or recataloged for recall as indicated by message FDR314. For Superscratch, this data set may have been uncataloged as indicated by message FDR314. For FDRCOPY MOVE, if the data set was cataloged, it has been successfully recataloged and/or uncataloged. A "status" of **BYPASSED** indicates that the scratch was not attempted for this data set, because of an earlier error, such as a catalog error, which is described by another message.

Action: If appropriate, scratch the data set manually. If desired, contact Innovation for assistance.

FDR333 CARD IMAGE--control statement image

Reason: Displays a control statement read from an ABR remote queue data set (DDname ABRREST for RESTORE TYPE=ABR, ABRARCH for RESTORE TYPE=ARC, ABRBKDQ for DUMP TYPE=ABR/AUTO, ABRARDQ for DUMP TYPE=ARC).

FDR334 FDR FAILED TO CATALOG comp DSN=dsname status

Reason: FDR allocated and restored the non-VSAM data set "dsname", but could not catalog it. If "comp" is printed, it will be **COMP=X'ffff00001111'** which contains the catalog return codes from registers R15 (ffff), R0 (0000), and R1 (1111). If the register 15 value is X'FF14', then the code was generated by FDR, and indicates either that the data set was currently cataloged on zero volumes or more than 255 volumes, or that the volume sequence number of the current piece of the data set was higher than 255; the R0 value is the number of volumes currently in the catalog entry (in hex). FDR supports cataloging of multi-volume data sets up to a maximum of 255 volumes. If "status" is **ALREADY CATALOG VOL=vvvvvv** then RECAT was not specified, and the output data set was already cataloged to a volume (vvvvvv) other than the output volume (or the input volume for a MOVE). For a multi-volume data set, vvvvvv is the volume in the catalog entry that corresponds to the volume sequence number of the portion of the data set that has just been processed. If "status" is **NAME IS AN ALIAS** then the data set name being cataloged was an alias of another data set in the catalog so FDR did not update it.

Action: Determine the cause and take appropriate action. If desired, contact Innovation for assistance.

**FDR335 ABR BYPASSED RESTORE ON EXISTING DSN=dsname
ANOTHER TASK COMPLETED RECALL OF DSN=dsname**

Reason: The first form of the message occurs when the PRESTAGE option was specified but the data set was found on the first output disk selected. The second form occurs when an auto-recall of an archived data set is requested but a previous request for the same data set is already recalling it.

100.03 CONTINUED . . .

- FDR336 DYNAMIC ALLOCATION ERROR COMP=cc, CODE=nnnn, INFO=iiii, DDNAME=ddname bypass**
Reason: FDR attempted to dynamically allocate (DYNALLOC) a disk volume or a data set (such as a disk or tape backup data set or the Archive Control File), but the allocation failed. "cc" is the return code in R15, "nnnn" is the dynamic allocation reason code and "iiii" is the dynamic allocation information code. For OS/390, these codes are documented in the IBM manual Authorized Assembler Services Guide. They can also be found in Appendix A of the ISPF online HELP. Frequently encountered "nnnn" values include:
- 0210** requested data set not available (e.g., another job had the Archive Control File with DISP=OLD).
 - 0218** TSO user did not have MOUNT privileges to mount an ABR tape.
 - 021C** invalid unit name used (e.g., ABR is trying to allocate a 3480 on a system without 3480s).
 - 0220** requested volume not available (e.g., another job is using a tape volume required by ABR).
 - 0484** the operator replied "NO" to Message IEF235D.
 - 04A8/04AC/04B0** allocation of a backup controlled by StorageTek's ExHPDM product failed. It may be that the ExHPDM address space is not running, or the ExHPDM subsystem name is not "SOV". If you use a different subsystem name, contact Innovation.
 - 1708** data set not found in the catalog (e.g., the name in the ABR option table for the Archive Control File does not exist in the catalog).
 - 7700** an error allocating a backup controlled by StorageTek's ExHPDM product. It may be that the ExHPDM address space is not running. If it is, check the ExHPDM log for errors.
- "bypass" may be **VOL=vvvvvv BYPASSED** which indicates that ABR bypassed processing the indicated volume because of the error or **DSN=dsname BYPASSED** if a data set was bypassed.
- Action:** This volume or data set operation will be bypassed. In some cases, the error can be avoided by setting the dynamic allocation flags in the FDR global option table (panel A.I.4.4, See [Section 90](#)). If necessary, contact Innovation for assistance.
- FDR337 ABR PRESTAGE--ALL DATA SETS ALREADY EXIST**
Reason: PRESTAGE was specified and all of the data sets to be restored were found to exist.
Action: ABR will terminate the operation with no data sets restored.
- FDR338 MAXIMUM BACKUP FILES EXCEEDED -- SPECIFY MAXBACKUP=NNNN TO OVERRIDE**
Reason: An ABR restore job was requested to restore data sets from more than the maximum separate backup data sets (default is 100 backup files). The restore job will fail with a U0628 ABEND.
Action: Specify MAXBACKUP=nnnn on the RESTORE statement or eliminate some of the requested data sets to reduce the number of backup files required.
- FDR339 BACKUP OF ARCHIVE FILE WAS BYPASSED**
Reason: This is a warning that ABR either was unable to backup the ARCHIVE control file, or was instructed not to, by the ARCBACKUP=NO operand.
- FDR340 RACF DEFINE/CHECK ERROR ON DSN=dsname COMP CODE=ccc**
Reason: FDR issued a RACF DEFINE or CHECK for the data set "dsname", but it failed with return code "ccc".
Action: Determine the cause.
- FDR341 VTOC CONVERT ROUTINE ERROR COMP=ccc**
Reason: A FDR function attempted to allocate a dummy data set with the DOS flag set in the VTOC in order to invoke the DOS VTOC conversion routine to recalculate the free space on a volume. This allocation failed with decimal return code "ccc"; for OS/390 these codes are documented under "allocate" in the IBM DFSMSdfp Diagnosis Reference manual. "ccc" of 172 indicates that your security system rejected the dummy data set (which begins with "FDRABR.Vvolser"). "ccc" of 176 or 180 indicates that a user-provided DADSM exit rejected the request. "ccc" of 192 indicates that the volume was SMS-managed (it will always fail on SMS-volumes).
Action: If this occurred on an Archive Backup, the archive of data sets from this volume is complete but the volume free space may not be accurate.
If it occurred on a full-volume restore to a different sized volume (e.g., 3390-2 to 3390-3), the restore is complete but the volume free space may not be accurate. If the volume had an active Indexed VTOC (VTOCIX), you must allocate a data set to correct the freespace and run ICKDSF BUILDIX to rebuild it.
The DOS flag will be left on, so the next allocation to the volume should correct the free space. Alternately, you can use COMPAKTOR to correct the free space

100.03 CONTINUED . . .

- FDR343 SMS CONSTRUCT ERROR VOL=vvvvvv construct REASON=reason**
Reason: FDRABR tried to get information about an SMS construct (e.g, storage class, management class, or base configuration) but the request failed.
 "vvvvvv" is the volume involved (if known), "construct" is STORGRP=storgrp for an error on a storage group construct, MANAGEMENT CLASS for an error reading all management class constructs, or BASE CONFIG for an error reading the SMS base configuration.
 "reason" is the error text:
SMS ERROR COMP=xxxx CODE=ccccc -- the IBM SMS interface failed the query with return code "xxxx" (usually 0008) and SMS reason code "ccccc". You can interpret the reason code by looking up IBM message IGDccccc in IBM message manuals.
STORAGE GROUP TYPE IS NON-POOL -- a VIO or DUMMY storage group name was specified.
STORGRP NOT ENABLED FOR MIGRATION -- for DUMP TYPE=ARC or TYPE=SCR, the storage group was not enabled for auto migration.
STORGRP NOT ENABLED FOR AUTO BACKUP -- for DUMP TYPE=ABR TYPE=DSF or TYPE=AUTO, the storage group was not enabled for auto backup.
STORGRP NOT ENABLED FOR AUTO DUMP -- for DUMP TYPE=FDR, the storage group was not enabled for auto dump.
STORGRP CONSTRUCT NOT FOUND -- the specified storage group name was not found by SMS.
STORGRP NOT FOUND OR NO VOLUMES -- The storage group is inactive or contains no active volumes. Determine which storage groups and volumes are available using ISMF or the operator command "D SMS,STORGRP(ALL),LISTVOL".
STORGRP MORE THAN 255 ACTIVE VOLS -- storage groups with over 255 volumes are not supported by MOUNT STORGRP=.
STORGRP CONTAINS NO ONLINE VOLUMES -- the specified storage group contains no volumes online to the operating system.
Action: ABR will bypass the volumes or storage group specified and continue. If necessary, correct the error and resubmit.
- FDR344 FDR FAILED TO FIND THE ABR HIGH LEVEL INDEX**
Reason: FDRABR issued a LOCATE (SVC 26) for the ABR high level index (by default, FDRABR), however, the alias could not be found.
Action: Use IDCAMS LISTCAT to list the master catalog and determine if an alias entry exists for the ABR index. Check the ABR high level index name specified in the FDR global options table and determine if the name matches the one in the master catalog. If cause cannot be determined, contact Innovation for assistance.

100.03 CONTINUED . . .

FDR345 FILTER ERROR REASON=reason -- ENTRY=entry
Reason: SELECT CATDSN=filter was used to select entries from the system catalogs, and an error occurred. "entry" indicates the catalog name or the catalog entry on which the error occurred. "reason" indicates the error:

- 1 -- NO CATALOGED ENTRIES FOUND** -- the filter did not select any entries from the catalogs.
- 2 -- CATLG ERROR COMP=cccc CODE=reason** -- The IBM catalog SVC returned return code "xxxx" with reason code "reason". These codes can be found under message IDC3009I in IBM message manuals.
- 3 -- CATLG NAME FPL ADDR MISSING** -- internal error.
- 4 -- CATLG VOLSER FPL ADDR MISSING** -- internal error.
- 5 -- CATLG VOLSER 0 OR MORE THAN 20** -- a catalog entry was selected that had either 0 volsers or more than 20 volsers in a CVOL (SYSCTLG) catalog.
- 6 -- MAXCARDS MUST BE INCREASED** -- since a SELECT statement is simulated for every data set selected from the catalog, the value of MAXCARDS= must be large enough to accommodate all selected data sets.
- 7 -- CATLG ENTYPE FPL ADDR MISSING** -- internal error.
- A -- NAME LAST CHARACTER HIGH VALUES** -- internal error.
- B -- VSAM COMPONENTS NOT FOUND** -- a ICF VSAM cluster was selected from the catalog but catalog entries for its components were not found.
- C -- VOL= CANNOT BE SPECIFIED** -- The VOL= operand can not be specified on a SELECT CATDSN= statement in a RESTORE operation.
- D -- CANNOT END IN A PERIOD** -- The filter cannot end in a period. [See Section 80](#) for filter rules.
- E -- INVALID GENERATION NUMBER** -- The filter ends in an invalid GDG relative generation number.
- F -- SEQUENCE ERROR IN CATALOG** -- A VSAM sequence error was encountered reading an ICF catalog, indicating a structural error in the catalog.
- G -- ALL ENTRIES EXCLUDED OR DUPS** -- all catalog entries selected by this filter were discarded because:
 - 1) they were excluded by a preceding EXCLUDE statement
 - 2) they were previously selected by a preceding SELECT CATDSN= statement
 - 3) they were migrated by HSM (volser-MIGRAT)
 - 4) they were archived by ABR and recataloged for Auto-recall.

Note: if you have volumes with VTOCs located above relative track 32767, CATDSN= will think they are cataloged for recall and will not select them. Innovation recommends that such VTOCs be relocated lower on the volume; COMPAKTOR can do this.
- H -- ERROR UPDATING CLONE X'vsamcode'** -- FDRCLONE had an error updating a catalog. "vsamcode" is the VSAM error code.
- J -- INSUFFICIENT STORAGE FOR CLONE** -- the above-the-line storage was insufficient for FDRCLONE. Increase the above-the-line region (use REGION=0M if you can) and resubmit.
- K -- ABR BACKUP NOT FOUND VOL=vvvvvv** -- the ABR backup file for a cloned data set was not found in the ABR catalog. The current backup of a selected data set is no longer recorded.
- L -- CANNOT CLONE IN MASTER CATALOG** - FDRCLONE cannot clone data sets cataloged in the master catalog.
- M-- CLONE ENTRY FOUND ON VOL=vvvvvv** - FDRCLONE was executed on a recovery system and the data set was found on disk, so it was not cloned.
- N -- CLONE BACKUP ALIAS DEFINE ERROR** -- FDRCLONE had an error defining an alias of a cloned data set.
- P-- CLONE DSN MISSING ON VOL=vvvvvv** -- when executing CLONE TYPE=DEFER, a data set selected for cloning is not currently on disk so FDRCLONE cannot determine its current backup.
- Q -- CLONE NO CURRENT BACKUP vvvvvv** -- a data set selected for cloning had no current ABR backup.
- R -- CLONE KEY RANGE NAME TOO LARGE** -- the data component name of a keyrange KSDS was larger than 38 characters. This should not occur.
- S -- CLONE OWNER CELL MISSING vvvvvv** -- the catalog entry for a data set to be cloned did not contain an owner cell. This should not occur.
- T -- CLONE MISSED UPDATE** -- due to an internal error, a data set was not marked as cloned.

Action: Serious errors will cause a U0502 ABEND. Others (such as reason 1) will simply be treated as if a SELECT statement did not select any data sets, unless SELTERR=NO was specified.

100.03 CONTINUED . . .

- FDR346** **FILTER SELECTED VOL=volser DSN=dsname**
FILTER SELECTED ALIAS=alias CAT=catalog TIME=hhmmss
FILTER SELECTED nnnnnnnn ENTRIES in sssss.ss SECONDS
Reason: A filter was used to select entries from the system catalogs. The first form of the message will be displayed only if PCATDSN= was used (instead of CATDSN=) and indicates each data set name selected from the catalog and the volser it is cataloged on.
The second form of the message will be displayed only if the PRTALIAS operand was specified, and indicates each alias which was selected from the master catalog by the filter and the name of the user catalog which was searched for it (even if no data sets were selected in that alias).
The third form of the message documents the number of entries selected by the filter, and the elapsed time in seconds required to complete the search.
- FDRCLONE CLONED nnnnnnnn FILES/CLUSTERS location**
FDRCLONE CLONED CLUSTER=clustername
FDRCLONE CLONED VOL=volser DSN=dsname
Reason: This form of the FDR346 message is displayed when you execute FDRCLONE. It displays the number of entries that were cloned, and identifies individual data sets and clusters which were cloned. "location" will be:
DIRECTLY IN THE CATALOG – TYPE=INPLACE
TO THE BACKUP FILE – TYPE=DEFER
FROM THE BACKUP FILE – TYPE=ACTIVATE
The second form of the message will be displayed only if the PRTALIAS operand was specified, and indicates each alias which was selected from the master catalog by the filter and the name of the user catalog which was searched for it (even if no data sets were selected in that alias).
The third form of the message documents the number of entries selected by the filter, and the elapsed time in seconds required to complete the search.
- FDR347** **SMS MGMTCLAS ERROR VOL=volser REASON=reason DSN=dsname**
Reason: SMSMANAGE=YES was specified, but an error occurred trying to use SMS management classes for ABR processing. "volser" identifies the volume being processed (if known) and "dsname" identifies the data set being processed for errors related to a particular data set. "reason" describes the error:
GETMAIN FAILURE -- MGMTCLAS DISABLED indicates that the GETMAIN for storage above the 16M line to hold the definitions of all active management classes failed. Management class processing will be disabled.
NVR/VVR TYPE UNKNOWN -- a VVDS record was not type N, Z, or Q. The VVDS may be damaged.
TYPE 23 NOT FOUND -- The type 23 cell was not found in a ICF VSAM VVR. The VVDS may be damaged.
TYPE 21 NOT FOUND -- The type 21 cell was not found in a ICF VSAM VVR. The VVDS may be damaged.
TYPE 22 NOT FOUND -- The type 22 subcell was not found in a non-VSAM NVR. The VVDS may be damaged.
TYPE 24 NOT FOUND -- The type 24 cell was not found in a non-VSAM NVR. The VVDS may be damaged.
MGMTCLAS INVALID -- The management class name in the NVR/VVR of a SMS-managed data set (or the default management class name) is not currently defined to SMS.
NO MGMTCLAS TABLE -- internal error.
NO VVDS ENTRY -- No NVR/VVR was found in the VVDS for a SMS-managed data set. The VVDS or the VTOC may be damaged, or creation of a SMS-managed data set was interrupted.
MULTI-VOL ERROR -- for a multi-volume data set, a LOCATE is done on other than the first volume to get the management class, and it failed.
Action: If DSN= is displayed, that data set will not be selected by ABR in this run. If possible, correct the error and rerun. If necessary, contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR348** **VOL=volser ALLOCATED SPACE IS aaa% -- THRESHOLD is ttt% - action**
VOL=volser ALLOCATED SPACE IS aaa% -- AFTER ARCHIVE/SCRATCH
VOL=volser BYPASSED
Reason: In Archive Backup or Superscratch step, volume threshold selection was requested via the THRESHOLD= operand and volume "volser" is being processed. In the first form of the message, "aaa" is the percentage of tracks currently allocated and "ttt" is the threshold percentage used by ABR on that volume (See Section 51 for details on the source of the threshold value).
 In the first form of the message, "action" will be:
 SELECTED DUE TO THRESHOLD if "aaa" is greater than "ttt".
 BYPASSED SPACE UNDER THRESHOLD if "aaa" is less than or equal to "ttt".
 The second form of the message will be printed after the Archive or Superscratch completes, giving the new percentage of tracks allocated after data sets were deleted.
 The third form of the message indicates that a LSPACE (SVC 78) issued to obtain the current allocation percentage, returned a non-zero return code. The volume will not be selected.
- FDR349** **COPY1 EXPIRES yyyy.ddd – COPY2 EXPIRES yyyy.ddd FOR DSN=dsname**
Reason: SMSEXPIRE=PRT was specified. This message is issued for every SMS-managed data set that is archived in this step, showing the calculated expiration dates for COPY 1 and COPY 2.
- FDR355** **A DUPLICATE NAME WAS SPECIFIED OR GENERATED --**
FDR BYPASSED DSN=dsname
Reason: NEWNAME=, NEWGROUP=, or NEWINDEX= was specified on one or more SELECT statements. However, when this was applied to the original data set name shown as "dsname", it resulted in an output data set name which duplicated the output name of another data set being restore or copied.
Action: Change the newname specification to generate unique names.
- FDR360** **CYL=cccc--cn cn cn...**
Reason: Internal diagnostic message optionally printed by the ABR full-volume restore processor showing each cylinder and head restored and the cycle number (cn) of the tape (0 thru 63) that it is restored from.
- FDR361** **TAPE WAS RESTORED DSN=dsname**
Reason: ABR successfully restored this tape on a full volume restore (RESTORE TYPE=FDR).
- FDR363** **OPERATOR CANCELLED RESTORE FOR TAPES**
Reason: The OPERATOR keyword was coded on the ABR RESTORE TYPE=FDR statement for a full-volume restore, but the operator replied "NO" to the FDRW25 message.
Action: The restore for this volume was cancelled.
- FDR364** **TAPE WAS BYPASSED BY OPERATOR DSN=dsname**
Reason: The OPERATOR keyword was coded on the ABR RESTORE TYPE=FDR statement for a full-volume restore, but the operator replied "BYPASS(nn)" to bypass the incremental backup "dsname".
Action: This incremental backup file will not be restored. Other incremental backups in the generation will be restored, which may cause some data sets to not be restored to their latest backup.
- FDR365** **TAPE WAS NOT FOUND DSN=dsname**
Reason: ABR attempted to locate this backup file in the ABR catalog. The catalog search failed. This entry may have been uncataloged.
Action: ABR will bypass this file. Use the TAPEDD operand on the SELECT statement if the data set name, volume serials and file number are known. Contact Innovation for assistance.
- FDR366** **CYL TRKS-0....5....ABCDEF0....5....ABCDEF -- TRACKS NOT FOUND ON TAPE**
Reason: The indicated tracks were expected to be on the backup tape, but they were not found. One detail line is printed for each cylinder for which tracks were not found. The cylinder number is printed in decimal. the track(s) that were not found are each indicated by an 'X'. The position of the 'X' indicates the number of the track: the first position after the hyphen in the heading line represents track 0, the next position is track 1, and so on.
Action: Review the listing from the backup run that created the tape. A possible reason for getting Message FDR366 on a restore is that Message FDR123 with REASON=2 appeared on the backup run, indicating that the entire cylinder was not dumped. On a data set restore, if message FDR155 does not appear, then the tracks not found were not needed, and the requested data sets have been correctly restored. On a full-volume restore, message FDR366 always indicates that not all of the tracks have been restored. If appropriate, contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR367 ABR TERMINATED RESTORE FOR NOT FINDING FDR TAPE**
Reason: On a full-volume restore, ABR could not locate the full-volume restore tape (cycle 00) in the ABR catalog. Message FDR365 identifies the backup data set.
Action: This volume restore is terminated. Contact Innovation for assistance.
- FDR368 TAPE WAS BYPASSED FOR LACK OF A VTOC DSN=*dsname***
Reason: On a full-volume restore, "*dsname*" was the cycle with which ABR was attempting to begin the restore, but ABR bypassed this tape since it was a manual backup (created with DUMP TYPE=DSF) and did not contain the VTOC.
Action: If the restore is from the ABR system, the next lowest cycle number in the current generation will be used to begin the restore. Any data sets on the bypassed backup will not be restored to their most current status.
- TAPE WAS BYPASSED BECAUSE ITS AN INSTANT BACKUP DSN=*dsname***
Reason: This form of the FDR368 message is issued during a full-volume restore, when ABR detects that a cycle to be restored, other than the first cycle read, was a backup on disk created by FDR InstantBackup which has not yet been moved to disk. ABR can restore from an InstantBackup on disk only if it is the first cycle being restored.
Action: Execute ABR to move that InstantBackup to tape, or use CYCLE= on the SELECT statement to select the cycle which is the InstantBackup.
- FDR370 OPEN ERROR -- TAPE BYPASSED DSN=*dsname***
Reason: A full-volume restore operation was requested from this tape. The open of the tape failed. ABR will bypass this tape unless it is the full volume (FDR) backup.
Action: Check for IBM OPEN messages in the joblog. Data sets that were on this incremental backup may not be restored to their latest status. Contact Innovation for assistance.
- FDR371 FORMAT 4 DSCB NOT FOUND ON RESTORE TAPES**
Reason: FDR or ABR full-volume restore did not find the Format 4 DSCB (first record of the VTOC) on any of the backup tapes. The restore may not be correct. This can occur if the pointer to the VTOC in the disk volume label was invalid at the time of the dump, or if errors occurred reading the backup records containing the VTOC
Action: Contact Innovation for assistance.
- FDR372 DISK FORMAT 4 READ ERROR**
Reason: An I/O error occurred reading the Format 4 DSCB (first record on the VTOC) of the output disk volume during a full-volume restore. The pointer to the next available alternate track on the volume, stored in the Format 4, if any, will be lost, but the volume is probably usable. A mini-dump will be printed, showing the I/O error.
Action: Run a COMPAKTOR MAP to check for VTOC errors. Contact Innovation for assistance.
- FDR373 UCB VTOC LOCATION WAS NOT UPDATED**
Reason: FDR failed to locate the disk label record on the backup tape (Cylinder 0 Track 0 Record 3) during a full-volume restore. The volume was restored, but the system pointers to the VTOC, VTOC INDEX, and VVDS were not updated. The output volume may not have a volume label.
Action: Use FDRDSF PRINT TYPE=DSF,DSN=VTOC to print the label track and VTOC to verify that the VOL1 volume label exists on CYL 0 TRK 0; if not, contact Innovation for assistance. If so, vary the volume offline and then remount it to update the system pointers.
- FDR374 BACKUP TAPE DID NOT CONTAIN THE ABR MODEL**
Reason: FDR failed to locate the ABR Model DSCB in the control records at the beginning of the first incremental backup read during an ABR full-volume recovery. A normal incremental backup always contains the ABR Model, so this is not normal (perhaps the ABR Model DSCB was excluded by an EXCLUDE statement).
Action: If the output volume contains an ABR Model DSCB and has the same serial as the volume being restored, the restore will terminate with a U0304 ABEND. If not, the restore will proceed; this message is issued once for each incremental that does not contain the ABR Model; some data sets may not be restored correctly.
- FDR375 ABR MODEL OR VTOC INDEX REWRITE/RENAME/ALLOC FAILED**
Reason: An error occurred when FDR attempted to rewrite or allocate the ABR Model DSCB, or rename the index VTOC or ABR model. The user may not be authorized to create, rename or update this data set.
Action: This volume may need to be reinitialized for ABR processing, or the indexed VTOC may need to be rebuilt with ICKDSF. Contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR376 MULTIPLE ABR MODELS FOUND ON BACKUP DSN=*dsname***
Reason: FDR found more than one ABR Model DSCB on the backup tape during a full-volume restore. The first one encountered is the one used. The disk volume serial number or the ABR model name may have been changed prior to the dump without using the ABR maintenance programs.
Action: SCRATCH the incorrect model using IEHPROGM specifying PURGE on the SCRATCH command. Use program FDRABRM with a REMODEL command ([See Section 50](#)) or the ABR ISPF dialogs ([See Section 90](#)) to reset the remaining ABR Model DSCB to the correct Generation and Cycle number. Contact Innovation for assistance.
- FDR377 DISK PACK HAS BEEN RELABELLED TO VOL=*vvvvvv***
Reason: Information message indicating that the volume serial of the disk being restored to has changed by a full-volume restore.
Action: If disk is shared by multiple systems, it should be varied offline and remounted to the other systems.
- FDR378 RENAME FAILED COMP=*cccc* CODE=*code* FROMDSN=*tempname***
NEWDSN=*permname* CLUSTER=*clusname*
Reason: On a FDRCOPY MOVE of an ICF VSAM cluster to the same name, the output cluster and its components were initially created with temporary names (an index level starting with T is inserted). After the input data set was deleted, FDRCOPY attempted to ALTER the temporary names back to the original cluster and component names but an error occurred. "cccc" is the return code and "code" is the reason code from ALTER; look up these codes under IBM message IDC3009I in the IBM message manuals.
"tempname" is the temporary cluster or component name, which FDRCOPY was trying to rename back to "permname". "clusname" is the original name of the cluster if the rename failed on a component.
Action: If a FDR378 is not issued for the cluster itself, it was renamed successfully and the cluster can be used under its original name.
If a FDR378 was issued for the cluster, the cluster was not renamed. It can be used under its temporary name "tempname", but you will probably want to check the error codes and ALTER the cluster to its original name if possible.
If FDR378 messages were issued for one or more components, those components were not renamed. The cluster is usable, but you may want to check the error codes and ALTER the component to its original name if possible.
If needed, contact Innovation for assistance.
- FDR390 DSF NO DATA SETS MET SELECTION CRITERIA**
Reason: No data sets were selected from the current disk volume. For an ABR step, this means that neither SELECT statements nor ABR automatic criteria (e.g., the update flag for Volume Backup) selected data sets on this volume.
Action: This may be normal if there were simply no data sets which met ABR requirements on the volume at this time. Otherwise check your SELECT statements and correct any errors.
- FDR391 OPERATOR REPLIED NO TO ARCHIVE ON UNEXPIRED DSN=*dsname***
Reason: ABR selected a data set for Archive Backup or Superscratch whose expiration date has not been reached. The operator replied 'NO' to the FDRW23 message.
Action: The data set will be bypassed. If the EXPD=NONE operand is specified on the DUMP statement, ABR will not perform this check.
- FDR392 function SELECTED DSN=*dsname* [(*clusname*)]**
Reason: Message showing which data sets were selected for backup. For ICF VSAM clusters, "dsname" is the component name and "clusname" is the cluster name.
- FDR393 USER PASSWORD EXIT FAILED ON DSN=*dsname* action**
Reason: User-written FDR data set security exit denied access to this data set.
Action: For a full volume operation (TYPE=FDR), the entire operation will be terminated and "action" will be **JOB TERMINATED**. On a data set operation (TYPE=DSF), the data set will be bypassed and "action" is blank. Check to be sure your data set security exit is properly denying access to the data set.
- FDR394 TAPE WAS BYPASSED BY DYN TAPE ERROR DSN=*dsname***
Reason: ABR attempted to dynamically allocate the backup file "dsname". The allocation failed. The FDR336 message gives the completion code.
Action: Determine the reason (such as the device was not available) and retry the operation.

100.03 CONTINUED . . .

- FDR395** **MAXIMUM BACKUP TRACKS PER FILE EXCEEDED (MAXBTRKS) - DEFERRING DATA SETS TO NEXT FILE VOL=volser**
Reason: During an ABR Archive Backup step, ABR selected data sets from volume "volser" totaling more tracks than the limit specified by the MAXBTRKS= operand (which defaults to 4096). In order to keep the size of individual Archive Backup files down and improve recall performance, ABR will stop selecting new data sets, complete Archive processing for the selected data sets, and reprocess this volume to select additional data sets. Multiple backup files will be created.
- FDR400** **program/function -- VER v.r/mmt -- INNOVATION DATA PROCESSING**
DATE=yyyy.ddd TIME=hh.mm.ss PAGE nnnn
Reason: General page header message for ABR programs other than FDRABR. See message FDR101 for details. This time is not reported by all programs.
- FDR401** **PARAM DATA -- * parm-field-data ***
Reason: Displays the program control information specified in the 'PARM=' field of the EXEC statement. This data will not be displayed if program is invoked under TSO.
- FDR402** **INVALID CONTINUATION**
Reason: The last control statement ended in a comma, but there was no next statement or it was not a proper continuation statement.
Action: Correct and resubmit job.
- FDR403** **REQUIRED OPERAND NOT SPECIFIED -- operand -- reason**
Reason: The operand indicated is required for the execution of the current statement. It must be specified; no defaults are available. "reason" may include some explanatory text.
Action: Refer to the documentation for the program you are executing in this manual. Correct and resubmit job.
- FDR404** **CONTROL STATEMENT LIMIT EXCEEDED**
Reason: The maximum number of control statements or operands has been exceeded. The message text may vary depending on the particular limit that was exceeded.
Action: Refer to the documentation for the program you are executing in this manual. You may be able to increase the limit. If not, you will need to change the job to reduce the number of statements or operands.
- FDR405** **MAXIMUM CONTINUATION COUNT OF nnnn EXCEEDED -- COMMAND FLUSHED**
Reason: One control statement used too many continuations statements.
Action: Combine multiple operands onto one line to reduce the number of continuations.
- FDR406** **CONFLICTING OPERANDS - operands**
Reason: The operands listed conflict with each other and are mutually exclusive. This message is normally followed by FDR407.
Action: Refer to the documentation in this manual for the program you are executing. Eliminate the conflict by removing one or more operands.
- FDR407** **type ERROR -- action**
Reason: An error was encountered during the processing of user supplied control statements. Always preceded by one or more messages which define and delimit the error(s).
 "type" will be either **PARAMETER FIELD** or **CONTROL STATEMENT**.
 The "action" is one of the following:
 JOB TERMINATED -- Processing will stop after the first error has been encountered.
 SKIPPING FOR COMMAND -- Processing will continue for all command statements found within the SYSIN data set.
 RE-ENTER COMMAND OR END -- Message for a user that has SYSIN data set assigned to a TSO terminal. Re-enter command in error or 'END' to complete the processing.
Action: Correct and resubmit job.
- FDR408** **NO CONTROL STATEMENTS WERE FOUND -- JOB TERMINATED**
Reason: SYSIN data set is empty OR contained only comment statements ("*" in column 1). The first control statement can also be supplied by the EXEC statement PARM= option. There are other forms of this message which indicate that statements required for a certain function were not found.
Action: Correct and resubmit job.

100.03 CONTINUED . . .

- FDR409 LENGTH OF op1 EXCEEDS THE LENGTH OF op2**
Reason: The number of significant characters specified for the "op1" operand exceeds the number of significant characters specified for the "op2" operand. They are interrelated and "op2" must be as long as or longer than "op1".
Action: Refer to the documentation in this manual for the program you are executing. Correct the operands specified and resubmit the control statement.
- FDR410 VOL=vvvvvv ASSUMED FOR DSN=dsname**
Reason: A command specified DSN= and did not specify VOL=. The data set "dsname" was a cataloged, single-volume data set. The program assumed that the volume where the data set was cataloged was the volume on which the user wanted the command to process the data set.
Action: If the assumption was correct, none. Otherwise, you may wish to reverse the effect of the command (e.g. by issuing RESET ARCHIVE if the original command was ARCHIVE) and then to reissue the original command specifying VOL=.
- FDR411 MISSING ABRINIT, MAINT, OR SELECT CONTROL STATEMENTS(S) -- JOB TERMINATED**
Reason: FDRABRM was executed with DEFAULT as the only control statement.
Action: Check your control statement listing, correct any errors found and resubmit the job.
- FDR413 DSN/DSG type ERROR -- description**
Reason: There is an error in the specification of the DSN= or DSG= operand. "description" may be:
MAX INDEX LEVEL EXCEEDED - a maximum of 22 leading periods are allowed,
NO SIGNIFICANT DATA - the leading periods were followed by a comma or blank
INVALID CHARACTERS - invalid characters in the name
INVALID GENERATION - a GDG generation number was specified, but it is invalid.
MAX INDEX LENGTH EXCEEDED - an index level exceeded 8 characters.
Action: Refer to the documentation in this manual for the program you are executing. Correct and resubmit.
- FDR414 IGNORED -- description -- DSN=dsname**
Reason: The device type returned for the data set was not that of a disk device or the data resides on multiple volumes. The processing requested for this data set is not performed.
Action: Correct and resubmit job.
- FDR415 ON VOLSER vvvvvv - DSCB tt-cccchhhrr-CHAINED TO DSCB tt-cccchhhrr - action**
Reason: An invalid DSCB chain was detected in the VTOC of volume "vvvvv". Information is displayed for the current DSCB and the invalid DSCB it chains to: "tt" is the DSCB type (1-7) and "cccchhhrr" is the record address (cylinder, track and record number) of the DSCB. "action" may be:
PROCESSING CONTINUES if the information in the DSCB is not critical for this operation
NOT FOUND IN INCORE VTOC if the chained DSCB is not found in a copy of the VTOC saved in storage. The operation is terminated.
Action: Even though the program may have been able to continue, this is a serious problem. Determine the data set(s) involved and attempt to correct the DSCB. You can use COMPAKTOR to help confirm and identify the error. If unable to determine the cause of the error, contact Innovation for assistance.
- FDR416 IMPROPER USE OF OPERAND -- description**
Reason: An operand was used improperly. "description" identifies the operand and the reason.
Action: Refer to the documentation in this manual for the program you are executing. Eliminate the bad operand or use it correctly.
- FDR417 USE OF RESTRICTED FEATURE -- description**
Reason: A feature has been disabled in the FDR Global Option Table. "description" identifies the feature.. Attempts to use a disabled feature results in the subcommand being marked in error.
Action: Eliminate the restricted feature. Rerun or restart the job.

100.03 CONTINUED . . .

- FDR418** **COMMAND PROCESSING DETECTED ERROR--*action***
Reason: An error was encountered during the processing of a statement. Always preceded by one or more messages that define the error. "action" will be one of:
RE-ENTER COMMAND OR END -- message for user that has SYSIN data set assigned to a TSO terminal. Re-enter command in error or 'END' to complete the processing.
SKIPPING FOR COMMAND -- processing will continue for all command statements found within the SYSIN data set.
REVERTING TO SYSIN -- the error occurred when reading from an alternate command input source. Processing will continue for commands in the SYSIN data set.
Action: Correct and resubmit job.
- FDR419** **CAUSED BY THE FOLLOWING COMMAND --**
Reason: One or more syntax errors on a remote queue were encountered during the execution of either a 'LIST' or 'REMOVE' subcommand by the remote queue utility, FDRABRUT. Messages defining the errors were listed prior to this message. The command that caused the errors is listed immediately following this message.
Action: Execute the REMOVE command specifying ARCHIVE or BACKUP, as required, and include the operand 'ERRORS'. This will remove any queue entries in error.
- FDR420** ***resource* PROCESSING BYPASSED - *reason***
Reason: Processing has been bypassed on "resource" (**DDNAME=ddname or VOL=volser**). "reason" may be:
VOLSER vvvvvv NOT INITIALIZED FOR ABR - volume does not contain an ABR Model DSCB
VOLSER vvvvvv HAS ARCHIVING DISABLED - The ABR Model DSCB does not have the flag enabling this volume for Archive Backups.
FDREPORT IS NOT AUTHORIZED - the RESERVE on the VTOC of the volume will not be done, but FDREPORT will continue.
VTOC ON CYLINDER 0/HEAD 0 - the VTOC on the volume is in a location used only for dummy VTOCs on volumes formatted for VM processing. It does not contain useful data so the volume is bypassed.
FORMAT 4 DSCB IS NOT THE FIRST ENTRY IN THE VTOC - the VTOC is corrupted.
THERE IS MORE THAN ONE (1) FORMAT 4 DSCB IN THE VTOC - the VTOC is corrupted.
INITIALIZATION REQUESTED BUT VOLUME IS ALREADY INITIALIZED - volume already contains an ABR Model DSCB. If you still want to reinitialize the volume, use the REMODEL statement.
VOLUME HAS NOT BEEN INITIALIZED - volume does not contain an ABR Model DSCB
VTOC CONTAINS DSCBS WITH NON-ZERO RESERVED FIELDS - during an attempt to initialize a volume for ABR, FDRABRM found that one or more Format 1 DSCBs (for the data sets currently on the volume) had non-zero values in the reserved bytes at displacement 103 and 104 in the DSCB, so the volume could not be used for ABR. This data may have been stored by another DASD management system. If you are sure it is no longer needed, specify the FORCE operand to force the initialization (bytes 103 and 104 will be zeroed). Contact Innovation for assistance.
ALLOCATION OF THE ABR MODEL FAILED - FDRABRM could not allocate the ABR Model DSCB.
Other reasons are possible.
VOLSER ALREADY IN USE - can occur if you provided DISKx DD statements and also requested dynamic allocation of disk volumes.
DUPLICATE DDNAMES - can occur if you provided DISKx DD statements and you used the same DISKx ddname more than once.
MULTIPLE UNITS REFERENCED - can occur if you provided DISKx DD statements and one such DD statement referenced multiple disk volumes.
MAXONLINE VALUE nnnn EXCEEDED - more than nnnn volumes were selected for processing in this step. You can increase the value of MAXONLINE on ISPF panel A.I.4.10, or break up the job into multiple steps.
VOLSER CONTAINS RESERVED CHARACTER 'X' - the volume serial of the disk contains a character which is not supported by MVS in volume serials.
Action: If the condition is one that can be corrected, correct it and resubmit the job. If needed, call Innovation technical support for assistance.

100.03 CONTINUED . . .

- FDR421 type LOCATE ERROR -- description -- DSN=dsname**
Reason: (1) A Catalog LOCATE SVC was issued requesting identification of the ICF component or SMS managed data set named by "dsname". The LOCATE either failed or returned a component type or cluster name description that is not currently handled.
 (2) The DIAGNOSEVVDS option was in effect, and an error was found in the VVDS.
 "description" can be:
RETURN CODE rc -- REASON IGGOCLaa-crs - the LOCATE failed. The error codes are documented in the IBM System Messages manual under message number IDC3009I.
ENTYPE - c(X'xx') - **entrytype** - the program encountered a type of component entry that it does not presently support.
NAMEDS - c(X'xx') - **entrytype** - the program encountered a type of cluster entry that it does not presently support.
NO MATCHING VVR -- ICFSOURCE=VVDS - A VSAM data set does not have a VVR in the VVDS.
NO MATCHING NVR -- SMSSOURCE=VVDS - A non-VSAM data set on an SMS-managed volume does not have an NVR in the VVDS.
MULTIPLE VVRS (nn) ON VOLUME vvvvvv - A VSAM component has nn VVRs in the VVDS.
MULTIPLE NVRS (nn) ON VOLUME vvvvvv - A non-VSAM data set has nn NVRS in the VVDS of SMS-managed volume vvvvvv.
 For any other description, an error exists in the catalog.
Note: This message only appears if the ICFERRPRT=YES or DIAGNOSEVVDS option has been specified within the FDR global option table or on the control statements.
Action: This is only a warning message. The program does not associate the component shown by DSN= with a cluster name, or attempts to recover from the VVDS error, and processing continues. If possible, correct the error and re-run.
 In some cases, the program may display a jobstream that you can use to correct the error. If you request the PUNCHDIAGNOS option, this jobstream will be written to a data set, specified by ddname DIAGPCH.
- FDR422 DDNAME=ddname REFERENCES DATA SET CURRENTLY IN USE--DSNAME=dsname**
Reason: The data set described by the message is being accessed from another task or address space.
Action: Rerun job when file is available for processing.
- FDR424 PROCESSING OF VOLUME SERIAL NUMBER vvvvvv COMPLETED--nnnn DATA SETS UPDATED**
Reason: Processing of volume "vvvvvv" completed. "nnnn" Format 1 DSCBs were changed (this includes the ABR Model DSCB if appropriate).
- FDR425 NOT REFERENCED DURING EXECUTION-- description**
Reason: Identifies a statement or operand for which no matching volumes or data sets were found during this execution.
Action: Verify all control statements are correct. This control statement did not affect the execution. Results may not be what was expected. Examine output carefully!
- FDR426 PROCESSING OF VOLUME vvvvvv COMPLETED**
Reason: The program has successfully completed processing for the disk volume "vvvvvv".
- FDR427 WARNING ARCHIVE FILE SEARCH TERMINATED AT DATE=yyyy.ddd**
Reason: SDATE or SDAYS was specified on the PRINT ARCHIVE control statement. ABR has terminated the search at this archival date. If data sets were archived prior to this date they will not be selected.
- FDR428 WARNING ONLY nn% OF ARCHIVE CONTROL FILE FREE--REORGANIZATION NEEDED**
Reason: During an Archive report, FDRABRP or FDREPORT detected that the Archive Control File is becoming full. The default value for the threshold of free space below which this message is produced is 10% of the file.
Action: Caution! Reorganize and/or enlarge the Archive control data set with FDRARCH ([Section 51](#)). You may change the threshold of free space below which this message is produced by using the REORG% operand of the PRINT ARCHIVE command.

100.03 CONTINUED . . .

- FDR429 WARNING nn DAYS SINCE VOLUME=vvvvvv PROCESSED BY ABR**
Reason: During a backup report, FDRABRP or FDREPORT detected that the disk volume "vvvvvv" has not been processed by ABR Volume Backups for "nn" days. The operand BKDAYS= specifies the period which invokes the message (default: 7 days). This message can occur for volumes which have been deleted from your system or which are no longer being backed up by ABR.
Action: If this message is unexpected, you should investigate why the volume is not included in your regular backups. You should probably arrange for a full-volume backup or incremental backup of the volume to be taken very soon. If your full-volume backups are taken at intervals of more than one week, increase the values of BKDAYS=. For volumes which are no longer required to be backed up by ABR, the PURGE BACKUP command of program FDRABRCM (See [Section 51](#)) can be used to delete the entries from the ABR catalog.
- FDR430 description--DSN=dsname**
Reason: User requested that an ICF VSAM cluster to be marked (RESET) for archive or backup processing. Message is issued in two (2) lines. The first line of the message identifies the cluster name (**ICF CLUSTER NAME**) and the second line of the message identifies the name of the component that will be marked as requested (**MARKING COMPONENT**).
- FDR433 MEMORY SPACE REQUEST FAILED BYTES REQUESTED--nnnnnnnn**
Reason: While processing ICF VSAM records from the Archive file, a table of size "nnnnnnn" (bytes, in hex) was required but the GETMAIN for it failed.
Action: Increase region size and rerun.
- FDR435 FILTER ERROR REASON=reason -- ENTRY=entry**
Reason: FDR\$CATP, the Catalog Read Service used by FDREPORT and other ABR programs, was invoked to extract information from the system catalogs, and an error occurred. "entry" indicates the catalog name or the catalog entry on which the error occurred. "reason" indicates the error:
1 – NO CATALOGED ENTRIES FOUND - the argument that was used failed to select any entries from the catalogs.
2 – CATALOG NOT FOUND BY OBTAIN - an OBTAIN SVC was issued to insure that a catalog was on a volume prior to allocation. The OBTAIN failed.
5 – USER REQUESTED BYPASS - the user requested that this catalog not be processed.
H – INSUFFICIENT STORAGE AVAILABLE ABOVE THE 16M LINE - all the ICF control blocks and buffers as well as program storage is acquired from above the 16M line and one of those GETMAINS failed.
K – COMMAND OUT OF SEQUENCE - FIRST COMMAND NOT MASK BUILD - the first command presented to the Catalog Read Service must be a Mask Build.
L – POINT FAILED CAT=catalog - a VSAM POINT command failed in the indicated catalog. This usually indicates that the catalog is damaged.
M – ALLOCATION FAILED - R15=xxxx - CODE=xxxx - CAT=catalog - dynamic allocation of the indicated cataloged failed with the return and reason codes displayed.
N – DEALLOCATION FAILED - R15=xxxx - CODE=xxxx - CAT=catalog - dynamic deallocation of the indicated cataloged failed with the return and reason codes displayed.
O – ccccc GENCB FAILED - R15=xxxx - CODE=xxxx - CAT=catalog - a VSAM GENCB was issued for the control block "cccc" but it failed with the return and reason codes displayed.
R – OPEN FAILED - R15=xxxx - CODE=xxxx - CAT=catalog - a VSAM OPEN was issued for the indicated catalog but it failed with the return and reason codes displayed.
X – ERROR READING CATALOG - R15=xxxx - CODE=xxxx - CAT=catalog - a VSAM READ was issued for the indicated catalog but it failed with the return and reason codes displayed.
Action: Serious errors will cause the requested service to be terminated. Others (such as reason 1) will simply be treated as if a SELECT statement did not select any data sets, unless SELTERR=NO was specified.
- FDR436 FAMS REQ BYPASSED - ALLOCATION FAILED - R15=xxxx - CODE=xxxx - DSN=dsn**
Reason: Dynamic allocation was requested for the HFS or PDSE data set indicated but it failed with the return and reason codes displayed.
Action: FAMS processing of the data set is bypassed.

100.03 CONTINUED . . .

- FDR440 WARNING BACKUP DATE *yyyy.ddd* IS GREATER THAN RUN DATE**
Reason: An entry on the ABR backup catalog contains a date greater than the current date known to the Operating System.
Action: Check if the date set for use by the Operating System is correct or the ABR run that created the entry was executed with the date set to other than the current date.
- FDR441 XDSNAME SYNTAX ERROR--*reason***
Reason: In FDREPORT, a XDSNAME= operand on a XSELECT or XEXCLUDE statement violated the syntax rules, as detailed by "reason"
Action: Check the syntax rules in [Sections 54](#) and [80](#) and correct.
- FDR442 LAST CONTROL STATEMENT ADDED TO *queue***
Reason: The last command control statement was added to a remote queue as identified by "queue".
- FDR443 TAPE XREF TABLE FULL -- ENTRY COUNT IS *nnnn* -- ENTRIES WILL BE BYPASSED**
Reason: A FDRABRP PRINT BACKUP or FDREPORT requesting information on ABR backups exceeded the tabling capacity of the program. The maximum entries currently available are printed. Tabling of additional ABR backups is bypassed.
Action: To increase the maximum number of tapes that can be maintained in the table, options MAXTSVX and MAXTCVX can be updated in the FDR global options table:
 MAXTSVX -- applies to FDRABRP PRINT BACKUP reports not specifying COMBINE, and to all reports from FDREPORT. The number may be from 10 to 300. The default is 100.
 MAXTCVX -- applies to FDRABRP PRINT BACKUP reports that specify COMBINE. The number may be from 100 to 32000. The default is 1000.
 The FDREPORT table size can also be overridden by the MAXTAPE=nnn operand on the DEFAULT statement.
- FDR444 DDNAME: *ddname* DSNAME: *dsname* action**
Reason: During an FDRTCOPY, this shows the catalog status of the output data set. "action" may be:
NOT ABR INPUT OR ABR=NO - output data set not cataloged because the input was not an ABR backup file or ABR=NO was specified
NOT CATALOGED THIS VOLSER - output data set not cataloged because the input data set was not cataloged to the input volume.
NOT CATALOGED - output data set not cataloged because the input data set was not cataloged at all.
CATALOGED - the output data set was successfully cataloged.
 There may be additional FDR444 messages documenting the characteristics of the input backup data set.
Action: Specify CAT=RECAT if you want to force the output data set to be cataloged.
- FDR445 RECORD LENGTH ERROR**
Reason: During a FDRABRP PRINT TVTOC or FDREPORT DATATYPE=TVTOC function, the length field in the FDR control record in the backup data set does not equal the length of the block actually read.
Action: Contact Innovation for assistance if problem reoccurs.
- FDR446 1ST RECORD NOT FDR HEADER**
Reason: During a FDRABRP PRINT TVTOC or FDREPORT DATATYPE=TVTOC function, the first record read from the backup data set was not the expected special FDR header record.
Action: A U0626 ABEND will be issued. Check that the input data set is a FDR/DSF/ABR backup and resubmit. Contact Innovation for assistance if the problem reoccurs.
- FDR447 DSF RECORD OUT OF SEQUENCE**
Reason: During a FDRABRP PRINT TVTOC or FDREPORT DATATYPE=TVTOC function, FDR control records were not in the correct order.
Action: A U0626 ABEND will be issued. Contact Innovation for assistance.
- FDR448 INVALID DEVICE TYPE CODE IN FDR HEADER**
Reason: During a FDRABRP PRINT TVTOC or FDREPORT DATATYPE=TVTOC function, the device type code recorded in the FDR header record on the backup data set was not one supported by this version of FDR/DSF/ABR.
Action: A U0626 ABEND will be issued. Insure that the version of FDRABRP you are executing is at or above the level that created the backup. Contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR449 REBLK VALUE SPECIFIED TOO SMALL FOR A *disktype* INCREASING TO *nnnnn***
Reason: When copying a backup data set with FDRTCOPY, REBLKS= specified a value less than the track size of the dumped disk + 28 bytes.
Action: FDRTCOPY has automatically increased the output blocksize to "nnnnnn".
- FDR450 TAPEIN *status* TO *ddname*. BLOCKS READ:*rrrrrr*, BLOCKS WRITTEN:*wwwwww***
Reason: FDRTCOPY has copied one file from TAPEIN to TAPEOUT (and optionally to TAPE2OUT as well "status" may be **SUCCESSFULLY COPIED** or **COPIED WITH ERRORS**; in the latter case, non-terminating errors occurred so check for other error messages. Blocks read ("rrrrrr") and written ("wwwwww") will be the same unless reblocking was requested. This message will be followed by Message FDR615 detailing the data sets read and written.
 An alternate form of the FDR450 message may be issued with more details of the TAPEIN and TAPEOUT/TAPE2OUT data sets.
- FDR451 LOCK GREATER THAN 32760 TO DISK OR REBLKS OR UNBLK ON SPLIT FORMAT...JOB TERMINATED**
Reason: FDRTCOPY detected that either:
 1) The output file is on disk, and the input file was created with FORMAT=NEW (causing blocks greater than 32760), or REBLKS= specified a value greater than 32760.
 2) The input file was created using FORMAT=SPLIT and REBLKS= or UNBLK was specified. A U0502 ABEND is issued.
Action:
 1) If the output file is on disk, and the backup being copied is not a dump of a 3380 or 3390, specify REBLKS=32760 or less, or assign the output to tape.
 2) If the input file was created with FORMAT=SPLIT, do not specify REBLKS= or UNBLK.
- FDR452 ARCHIVING BYPASSED--*reason***
Reason: User requested that the data set identified in the FDR456 message that follows be marked for archive processing, but it was not marked because of the reason listed.
- FDR453 INDICATOR *action* FOR *type***
Reason: The data set identified in the FDR456 message that follows was SET or RESET for remote queue ARCHIVE or BACKUP dump processing.
- FDR454 INDICATOR *status* FOR *type***
Reason: User requested that the data set identified in the FDR456 message that follows be SET or RESET for remote queue ARCHIVE or BACKUP processing, but the indicator was ALREADY SET or NOT SET.
Action: Check the control statements to determine if they were properly specified. Correct as required.
- FDR455 *description***
Reason: Internal error in FDRTCOPY processing of REBLK or UNBLK reblocking.
Action: A U0659 ABEND is issued. Contact Innovation for assistance after obtaining a storage dump.
- FDR456 VOL/VOLG=*vvvvvv*--DSN/DSG=*dsname*--DDNAME=*ddname***
Reason: Identifies a volume, data set, and/or DDname, depending on circumstances. Will always be preceded by one of the following messages: FDR452, FDR453, FDR454, or FDR459.
Action: If preceded by FDR454, see action for that message, otherwise none.
- FDR457 ENQ FAILED--RC=*nn*--QNAME=*qname*--RNAME=*rname***
Reason: An attempt was made to RESERVE or ENQ on the resource "qname/rname" but it failed with ENQ return code "nn".
Action: Usually followed by a U0659 ABEND. Call Innovation technical support for assistance after obtaining a storage dump.
- FDR458 DEQ FAILED-- RC=*nn*--QNAME=*qname*--RNAME=*rname***
Reason: An attempt was made to DEQ the resource "qname/rname" but it failed with DEQ return code "nn".
Action: Usually followed by a U0659 ABEND. Call Innovation technical support for assistance after obtaining a storage dump.
- FDR459 *function*--DENIED BY *reason***
Reason: "function" was requested by the user on the data set identified by the FDR456 message that follows but the request was denied because of the security "reason" shown.

100.03 CONTINUED . . .

- FDR460 WARNING -- VOLUME CONTAINS MULTIPLE ABR MODEL DSCBS -- OTHER ABR FUNCTIONS MAY FAIL**
Reason: A VTOC maintenance request detected multiple ABR Model DSCBs on the disk volume. Any data set having the following format is considered to be an ABR Model DSCB: 'FDRABR.Vvolser' (note that the FDRABR high-level index may be changed in your installation). Multiple ABR Model DSCBs are usually caused by restoring the model DSCB from another volume.
Action: Delete the DSCB in error. A rerun of the job is not required.
- FDR461 ABR MODEL FOR VOLUME=vvvvvv NOT AVAILABLE AT THIS TIME USER MUST SPECIFY 'MAXGEN=' OPERAND**
Reason: An OBTAIN was issued to retrieve information from the ABR Model DSCB, but failed for some reason. The volume selected may not be online.
Action: The user must use operand 'MAXGEN=' to replace the information normally taken from the ABR Model DSCB, or supply a DD statement to ensure that the volume is mounted.
- FDR462 CURRENT GEN=gggg/CYCLE=cc ON THE ABR CATALOG IS IN CONFLICT WITH CURRENT GEN=gggg/CYCLE=cc ON THE ABR MODEL FOR VOLUME(vvvvvv) ALL BACKUP ENTRIES FOR THIS VOLUME BYPASSED.**
Reason: The current generation and cycle number recorded in the ABR Model DSCB on the volume "vvvvv", does not match the generation and cycle number of the entry with the most current date in the ABR backup catalog.
Action: Processing of the volume is bypassed. If you wish to process the volume anyways, use operand 'CURRINFO=MODEL/CATLG' indicating the correct value.
- FDR463 VOLSER=vvvvvv NOT INITIALIZED FOR ABR -- ARCHIVE CONTROL FILE MUST BE ON AN INITIALIZED VOLUME**
Reason: The user attempted to create an Archive Control File with either the FORMAT or RESTORE command. ABR requires that the Archive Control File reside on a disk volume that contains an ABR Model DSCB. FDRARCH will prohibit the creation of an ARCHIVE control file on non-initialized disk volumes.
Action: Initialize the volume for ABR or choose an ABR-initialized volume and rerun the job.
- FDR464 DDNAME=ddname REFERENCES AN UNSUPPORTED DISK DEVICE**
Reason: The disk device allocated to "ddname" is not in the ABR supported device table.
Action: Call Innovation technical support for assistance.
- FDR465 UNABLE TO OPEN DDNAME=ddname--reason**
Reason: An error occurred trying to open "ddname", for the "reason" displayed.
Action: Check the joblog for IBM OPEN messages. In some cases, the DDname is not required so the message is simply informational. If the DDname is required for the operation being attempted, correct and resubmit the job.
- FDR466 statement CANNOT BE PROCESSED**
Reason: A statement which requires queuing to a remote queue data set cannot be processed because the remote queue data set is not present. Always preceded by FDR465.
Action: Include the required DDname and resubmit the job including only those commands which failed.
- FDR467 UNABLE TO LOCATE HIGH INDEX OF iiiiii ON CATALOG--JOB TERMINATED**
Reason: The catalog manager has returned a condition code other than zero when attempting to locate the high level index "iiiiiii".
Action: It is possible that the ABR catalog has not been properly created. [See Section 90](#). If needed, call Innovation for assistance.
- FDR471 DDNAME=ddname I/O ERROR--SYNAD=synadtext**
Reason: A permanent I/O error was detected on the data set referenced by "ddname". The IBM macro SYNADef was invoked to format the cause of the error; the text returned is "synadtext".
Action: Examine the SYNADAF text to determine the cause of the error. Check the joblog for IBM I/O error messages. For additional assistance, please call Innovation.
- FDR472 DDNAME=ddname DOES NOT REFERENCE AN ARCHIVE CONTROL FILE-- DSN=dsname**
Reason: 1) "ddname" references an Archive Control File which has not been initialized by FDRARCH.
2) The Archive Control File has been corrupted.
Action: If the file was corrupted, you must restore it from the most recent backup. Otherwise, initialize it. Please call Innovation technical support if this problem occurs.

100.03 CONTINUED . . .

- FDR473 ABR INITIALIZATION OF VOLUME vvvvvv INCOMPLETE--
RENAME FAILED -- RETURN CODE rc -- STATUS ss**
Reason: During ABR initialization FDRABRM may rename the temporary data set allocated by the user to change it into the ABR Model DSCB. The rename failed. All other functions have been completed, including marking the Format 1 DSCBs for ABR processing. This problem can be caused by another task renaming or scratching the data set between the DEQ of the VTOC and the rename.
Action: If another task renamed or scratched the data set, rerun the FDRABRM job specifying FORCE on the appropriate ABRINIT control statement or the DEFAULT control statement. If the reason cannot be determined, call Innovation technical support for assistance.
- FDR474 DDNAME=ddname REFERENCES DATA SET WITH NO SPACE**
Reason: The data set was allocated with SPACE=(CYL,0) or (TRK,0) or (blocksize,0).
Action: Scratch the data set, re-allocate with sufficient space and resubmit job.
- FDR475 DDNAME=ddname REFERENCES DSNAME=dsname--NEED LEVEL PREFIX prefix**
Reason: 1) For program FDRARCH, the user coded a non-standard data set name for the Archive Control. The name must contain an index level of "ARCHIVE".
2) For program FDRABRUT, the user coded a data set name for a remote queue data set that did not contain the standard index (ABRARDQ, ABRARCH, ABRBKDQ, or ABRREST).
Action: Correct the dsname or (for FDRARCH only) specify DISABLE=DSNCK.
- FDR476 DDNAME=ddname REFERENCES A DEVICE TYPE OTHER THAN DISK**
Reason: "ddname" can only be allocated to a disk device. Check the unit specification in the JCL for errors.
Action: Correct and resubmit.
- FDR477 DDNAME=ddname REFERENCES A PREVIOUSLY FORMATTED ARCHIVE CONTROL FILE
'ERASE' MUST BE SPECIFIED TO CONTINUE**
Reason: Caution! You have attempted to FORMAT or RESTORE to an initialized ARCHIVE control file!
Action: If you want to clear the present file you must code ENABLE=ERASE on the FORMAT or RESTORE control statement.
- FDR478 nn RECORDS function**
Reason: Indicates that FDRARCH has performed the requested 'function' on 'nn' records. If NO RECORDS is printed, the function was not performed, probably due to a control statement specification.
Action: If NO RECORDS, correct the error and submit again. If NO RECORDS SELECTED FOR MAINTENANCE is printed, FDRARCH was invoked from FDRTCOPY when copying an archive backup, to update the Archive Control File to point to the output file(s), but no records were found on the Archive Control File which pointed to the input backup files being copied. This could occur if the input files had expired and being purged from the ACF, or if the input files had been previously copied so that the ACF no longer pointed to them.
- FDR479 nn RECORDS MERGED -- from**
Reason: Indicates that the number of records read by FDRARCH during a MERGE function. "from" is either **FROM DSN=dsname** to document records read from one MERGE input or **TOTAL OF ALL INPUT FILES** to show the total records read. If **NO RECORDS** is printed, no merge was performed, probably due to a control statement or JCL specification.
Action: If NO RECORDS, correct the error and submit again.
- FDR480 DDNAME=ddname REFERENCES AN EMPTY DATA SET**
Reason: The data set referenced by "ddname" from which the Archive Control File was to be restored contained no data.
Action: Verify that the tape specified is correct. If you are using other than SL tapes, check the LABEL= parameter. Correct the JCL and resubmit the job.
- FDR481 DDNAME=ddname DOES NOT REFERENCE
A BACKUP COPY OF AN ARCHIVE CONTROL FILE**
Reason: The data set referenced by "ddname" does not contain a backup copy of the Archive Control File in FDRARCH format.
Action: Check JCL statements and correct the DD statement specified in the message. FDR/DSF copies of the ARCHIVE control file cannot be restored by FDRARCH. They can be restored with the RESTORE TYPE=ABR statement of FDRABR or the RESTORE TYPE=DSF statement of FDRDSF.

100.03 CONTINUED . . .

- FDR482 DDNAME=ddname DOES NOT REFERENCE A TAPE DEVICE--REQUIRED FOR function**
Reason: "ddname" must point to a tape for "function".
Action: Correct JCL statement in error.
- FDR484 INTERNAL LOGIC ERROR--JOB TERMINATED**
Reason: The program has encountered an illogical condition.
Action: Call Innovation for technical assistance after obtaining a storage dump.
- FDR485 SORT FAILURE HAS OCCURRED RC=nnn--function CANCELLED**
Reason: Your installation's SORT product has returned a condition code other than zero. SORT error messages may be printed on Ddname SYSOUT or the job's joblog. Check your SORT documentation to interpret the error
Action: The processing of "function" has been cancelled. Other commands may process successfully. If you cannot correct the problem from this information, call Innovation for technical assistance.
- FDR486 VTOC READ ERROR--VOLSER=vvvvvv--SEEK ADDRESS cccchhhh -- action**
Reason: An I/O error was encountered while reading the VTOC of a disk volume. The volume serial "vvvvvv" and track address "cccchhhh" are displayed.
Action: Depending on user options, processing may continue bypassing remainder of VTOC or terminate with a storage dump. Call Innovation for technical assistance after obtaining a storage dump as a serious problem may exist. A COMPAKTOR MAP may help to diagnose the problem.
- FDR487 VTOC WRITE ERROR--VOLSER=vvvvvv--SEEK ADDRESS cccchhhh -- action**
Reason: An I/O error occurred while rewriting the VTOC of disk volume "vvvvvv". The volume serial "vvvvvv" and track address "cccchhhh" are displayed. Depending on user options, processing may continue bypassing the remainder of the VTOC or terminate with a storage dump. A COMPAKTOR MAP may help to diagnose the problem.
Action: Depending on user options, processing may continue bypassing the remainder of the VTOC or terminate with a storage dump. A COMPAKTOR MAP may help to diagnose the problem. Call Innovation for technical assistance after obtaining a storage dump as a serious problem may exist.
- FDR488 VTOC OUT OF SEQUENCE--VOLSER vvvvvv--SEEK ADDRESS cccchhhh--action**
Reason: ABR VTOC read routines check the CCHHR portion of the count field for continuity. During this check, one of the records was found to be in error. This message is usually followed by a diagnostic mini-dump. Depending on user options, processing may continue, bypassing the remainder of the VTOC, or may terminate with a storage dump.
Action: Check the mini-dump for the record in error. Check the volume for hardware problems. The VTOC may be corrupted; a COMPAKTOR MAP may help to diagnose the problem. If all else fails, call Innovation technical support for assistance.
- FDR489 NO prefix PREFIXED DDNAMES PRESENT OR SUITABLE FOR PROCESSING**
Reason: Depending on the function, "prefix" will be DISK or TAPE. No such DD statements were found in the step, or those that were present were rejected for other reasons, documented by other messages.
Action: Correct any error conditions or include the appropriate prefixed DD statements and resubmit the job.
- FDR490 PROCESSING TERMINATED WITHOUT FILE/VTOC UPDATE/LIST**
Reason: The processing program was unable to complete any of the functions requested. Always preceded by one or more error messages.
Action: Correct error conditions noted and resubmit the job.
- FDR491 function FUNCTION STARTED--hh.mm.ss**
Reason: Identifies type of function and time the function started.
- FDR492 function FUNCTION ENDED--hh.mm.ss CONDITION CODE-nnnn**
Reason: Identifies type of function and time the function ended. The return code is printed if non-zero.
- FDR493 function--BYPASSED--PRIOR FUNCTION TERMINATED WITH KEYWORD OR COMMAND DETECTED ERROR**
Reason: A previous command upon which this function may be dependent encountered serious errors.
Action: Correct the condition and reexecute.
- FDR494 WARNING INFORMATION REQUESTED NOT AVAILABLE FOR PROCESSING**
Reason: The information the user requested could not be located for print processing.
Action: Verify all selection control statements are correct.

100.03 CONTINUED . . .

- FDR495 NO DATA SETS MATCHED SELECTION CRITERIA--VOLSER=vvvvvv**
Reason: The selection criteria specified did not cause any data sets to be selected from the disk volume "vvvvvv" or from all disk volumes selected.
- FDR496 MODULE *modname* NOT USABLE WITH SYNC LEVEL *nnn* PROGRAMS - EXECUTION TERMINATED**
Reason: The main program determined that a subprogram was not at a compatible level, probably because different versions of the FDR program library are being used.
Action: Contact Innovation for assistance.
- FDR497 FDRTSEL *function* ENDED. ffff FILES PROCESSED OUT OF sssss SELECTED. *reason***
Reason: FDRTSEL ended without processing all selected files; "function" will be **PROCESSING** or **SIMULATION**. Out of "sssss" files selected for processed, "ffff" were actually processed. "reason" will be **STOP REQUESTED BY OPERATOR** or **MAXFILES LIMIT REACHED**.
- FDR498 *program (version)* PROCESSING COMPLETED WITH ERRORS**
Reason: "program" completed the required processing but non-terminating errors were encountered.
Action: Check the output, correct and rerun or restart as appropriate.
- FDR499 *program (version)* PROCESSING COMPLETED**
Reason: "program" completed the required processing successfully.
- FDR501 REMOVED FROM *ddname-- statement***
Reason: The user requested that "statement" be removed from the remote queue data set "ddname".
- FDR502 SIMULATED REORGANIZATION--*function status***
Reason: The user requested a simulated reorganization of the ARCHIVE control file. A "function" usually done during a REORG was BYPASSED or SIMULATED.
- FDR503 *reason* ARCHIVE ENTRY DROPPED--DSNAME=*dsname***
Reason: The entry for data set name "dsname" has been dropped from the Archive Control File.
"reason" can be:

RESTORED	the data set has been restored to disk.
EXPIRED	all Archive Backup files for this data set have expired.
DELETED	the data set has been marked for deletion.
ICF COMP	this is a component of a dropped cluster
ICF ERR	this is a component without any related cluster entry
IFNOTAUT	the data set is not marked for auto recall.
IFNOTCAT	the data set is not currently cataloged for auto recall.
MAXGENER	the data set exceeded the MAXGENERATIONS= value.
MAXOCCUR	the data set exceeded the MAXOCCURANCES= value.
- FDR504 DDNAME=*ddname* REFERENCES SUCCESSFULLY *function* ARCHIVE CNTL FILE**
Reason: The Archive Control File indicated by "ddname" has been **RESTORED, REORGANIZED, or MERGED** from other files.
- FDR505 DDNAME=*ddname* REFERENCES ARCHIVE CNTL FILE--DEVICE=*ttttt*--VOLSER=vvvvvv--DSN=*dsname***
Reason: Describes the device type, disk volume serial number, and the first 26 bytes of the data set name of the Archive Control File being processed.
- FDR506 SPACE AVAILABLE FOR *nnnnnn* ADDITIONAL ARCHIVED DATA SETS**
Reason: Lists the remaining capacity of the Archive Control File identified by messages FDR504 and FDR505.
- FDR507 DDNAME=*ddname* REFERENCES A BACKUP COPY OF DSNAME=*dsname***
Reason: A REORG or DUMP of an Archive Control File was requested by the user. The dump of the Archive Control File "dsname" was successfully completed to "ddname".
- FDR508 RECFM=F,BLKSIZE=*nnnnn***
Reason: Follows message FDR507 and documents the record format and blocksize of the backup file.

100.03 CONTINUED . . .

- FDR509 REJECTED--reason--VOL=vvvvvvv ADATE=yyyy.ddd DSNAME=dsname**
Reason: User has requested that one or more data set entries within an Archive Control File be marked for special processing or have a special processing indicator removed. The request was rejected for "dsname" that was ARCHIVED on date "yyyy.ddd" from disk volume "vvvvvv".
 "reason" can be:
ALREADY SELECTEDUser requested a data set entry be marked for Archive restore. It was already so marked.
EXPIRED TAPE User requested a data set entry be marked for Archive restore but the tape containing the data set has expired.
ALREADY DELETEDUser requested a data set entry be marked for deletion from the Archive Control File. The data set entry is already marked for deletion.
ICF TABLE FULL the table used for associating ICF VSAM clusters and components has overflowed.
VOLSER TABLE FULLthe table used to store backup volume serials has overflowed.
Action: Correct any obvious errors. Contact Innovation for assistance.
- FDR510 action--VOL=vvvvvvv ADATE=yyyy.ddd DSNAME=dsname**
Reason: User has requested that one or more data set entries within an Archive Control File be marked for special processing or have a special processing indicator removed. The requested "action" was successfully performed for the data set entry "dsname" that was Archived from disk volume "vvvvvv" on date "yyyy.ddd". This message may also appear during a REORG operation. "action" may be:
SELECTED FOR ARCHIVE RESTOREUser requested a data set entry be marked for Archive restore.
EXCLUDED FROM ARCHIVE RESTOREUser requested a data set be marked for exclusion from Archive restore.
SELECTED FOR RECORD DELETIONUser requested a data set entry be marked for deletion from the Archive Control File.
DELETION FROM ARCHIVE RESETUser requested a data set entry be reset for deletion from the Archive Control File.
UNCATALOGED FROM AUTO RECALLUser requested a reorganization of the Archive Control File. The specified data set entry which was being dropped from the file, had been marked for Auto Recall. Therefore the program uncataloged the data set.
- FDR511 RECALL DATA SET function status--DSNAME=dsname**
Reason: An FDRARCH function may create, update, or uncatalog the catalog entry for a data set archived for auto-recall. An error occurred during this process so the catalog entry for data set "dsname" was not disturbed. "function" will be **LOCATE, CATALOG, RECATLG, UNCATALOG** or **CVAFDIR**. "status" will be **FAILED** for an actual failure or **IGNORED** when **FDRARCH** could not do the catalog update (an additional reason is displayed). This message is always followed by message FDR516 or FDR674 to document the error when "status" is **FAILED**.
Action: See message FDR516 or FDR674 if "status" is **FAILED**.
- FDR512 WARNING--IMBEDDED INACTIVE BLOCK--**
ENTRY COUNT=nnn--RBN=rrr
Reason: The DUMP function of the FDRARCH utility encountered an imbedded block within an Archive Control File that is not marked as containing data. The number of entries within the block is printed as nnn. The relative block number within the Archive Control File is printed as rrr.
Action: This is a serious error. Call Innovation for technical support.
- FDR513 type TABLE CONTAINS nnnnnnnn ENTRIES**
Reason: A reorganization of the Archive Control File is being performed and obsolete Archive Backup data sets are to be uncataloged (UNCAT=YES). FDRARCH builds tables of all Archive Backups it finds in the Archive Control File. The resulting table contains "nnnnnnnn" entries. "type" may be **TAPE** for the full table of Archive Backup files, or **VOLSER** for the table of unique backup volume serials.

100.03 CONTINUED . . .

- FDR514 ARCHIVE TAPE *function* FAILED--DSNAME=*dsname***
Reason: A reorganization of the Archive Control File is being performed and obsolete Archive Backup data sets are to be uncataloged (UNCAT=YES) but a catalog action failed on the Archive Backup "*dsname*". "*function*" may be:
UNCATALOG - uncatalog of an obsolete Archive Backup failed
SCRATCH - scratch of an obsolete Archive Backup data set on disk failed
CATALOG - an Archive Backup tape data set that is no longer required had been uncataloged, but then it became necessary to catalog another Archive Backup tape data set on the same volume (Cataloging of Archive tape data sets is done only to provide an entry in the catalog for a given tape volume serial and is designed for those installations tape management catalog control retention).
Action: See message FDR516, FDR526 or FDR674 for details on the error.
- FDR515 ARCHIVE TAPE *action* VOL=*volser*--DSNAME=*dsname***
Reason: A reorganization of the Archive Control File is being performed and obsolete Archive Backup data sets are to be uncataloged (UNCAT=YES). "*action*" may be:
UNCATALOGED - an obsolete Archive Backup data set was uncataloged.
SCRATCHED - an obsolete Archive Backup data set on disk was scratched
CATALOGED - an Archive Backup tape data set that is no longer required had been uncataloged, but then it became necessary to catalog another Archive tape data set on the same volume (see message FDR514 for explanation).
- FDR516 CAMLST REGISTERS -- R0=xxxxxxx R1=xxxxxxx R15=xxxxxxx**
Reason: A CAMLST request for a CATALOG operation or a SCRATCH failed. The contents of Registers 0, 1 and 15 are displayed in hex; R15 will contain the return code but the other registers may have additional diagnostic data. The type of CAMLST function is identified in a preceding message. The return codes are documented in the IBM OS/390 manual DFSMSdfp Advanced Services.
Action: Check the return codes and take corrective action or call Innovation for assistance.
- FDR517 DSN--*dsname* FNUM--*nnn***
Reason: Identifies the data set name and tape file number (0 for disk) of an Archive Backup found in the Archive Control File during reorganization. Followed by messages FDR518 and FDR519 which further describe the tape table entry.
- FDR518 USE--*nnnnn* CATALOGED--*ccc* DEVICE--*ddddd* EXPDATE--*yyyy.ddd***
Reason: Follows message FDR517, and indicates the number of Archive Control File entries which reference the data set ("*nnnnn*", the catalog status of the file (YES or NO), the device type of the backup file, and its expiration date.
- FDR519 VOLS--*vvvvv* ...**
Reason: Follows message FDR518, and indicates the volume serial numbers of the tape or disk backup data set.
- FDR520 PRINTING FROM DEVICE TYPE *type* TO OUTPUT DDNAME *ddname***
Reason: An FDRDSF PRINT function was requested. "*type*" is the device type of the disk being printed, and "*ddname*" is the DDname to which the print output is directed.
- FDR521 CYLINDER *cccc* TRACK *ttt* -- RECORD ZERO *cccchhhrrrklll* -- DATA *xxxxxxxxxxxxxxxx***
Reason: An FDRDSF PRINT function was requested. For each track printed, this header identifies the track. "*cccc*" and "*hhhh*" are the cylinder and track number in decimal. "*cccchhhrrrklll*" is the count field of Record 0 (R0) on that track, in hex and "*xxxxxxxxxxxxxxxx*" is the 8-bytes of data from that Record 0, in hex.
- FDR522 COUNT FIELD *cccchhhrrrklll* -- eof**
Reason: An FDRDSF PRINT function was requested. For each record printed, this message displays the count field, in hex, of the record. "*cccchhh*" is the cylinder and track of the record, "*rr*" is the record number on that track, "*kk*" is the key length and "*lll*" is the data length. If the data length is zero, indicating an End-of-File record, "*eof*" will be **END-OF-FILE**, otherwise it is blank. The message is followed by displays of the key field (if "*kk*" is non-zero) and data field (if "*lll*" is non-zero) in dump format (hex and EBCDIC).

100.03 CONTINUED . . .

- FDR526 REGISTER 15=xxxxxxx--nnn VOLUMES REFERENCED.**
VOLSER=vvvvvv -- DEVICE TYPE=ttttt RETURN CODE=rrrr
Reason: Follows message FDR514 for SCRATCH failures. "xxxxxxx" is the SCRATCH return code in register 15 and "nnn" is the number of volumes on which the data set resides. The second line of the message is repeated for each volume, giving the SCRATCH status code for the volume. Code 7 indicates that SCRATCH failed because the data set was in use. Code 2 or 8 indicates a security violation.
Action: Check the codes and take corrective action or call Innovation for assistance.
- FDR530 member NOT FOUND--ddname=dsname**
Reason: Program issued a BLDL for the module "member" in the library "dsname" specified on the DD statement "ddname", but the module was not found. The module is required.
Action: Verify that the data set name specified on the DD statement was the correct one. If incorrect, change and resubmit the job. If the data set name is correct, check to see if all the installation steps required were done ([see Section 90](#)). Contact Innovation for technical assistance.
- FDR531 MODULE member CONTAINS NO TEXT RECORDS--ddname=dsname**
Reason: Program attempted to read module "member" in the library "dsname" specified on the DD statement "ddname", but no text records were found. The module is required.
Action: See message FDR530.
- FDR532 MODULE member I/O ERROR READING--ddname=dsname**
Reason: Program attempted to read module "member" in the library "dsname" specified on the DD statement "ddname", but an I/O error occurred. The module is required.
Action: Other messages may be displayed to further define the I/O error. Check the joblog for IBM error messages. Contact Innovation for technical assistance.
- FDR533 MODULE mmmmmmmm I/O ERROR WRITING--ddname=dsname**
Reason: Program attempted to update module "mmmmmmmm" in the library "dsname" specified on the DD statement "ddname", but an I/O error occurred. The update was not completed, but the module might be partially updated. This is a serious error which may result in the load module library becoming unusable.
Action: Other messages may be displayed to further define the I/O error. Check the joblog for IBM error messages. Contact Innovation for technical assistance.
- FDR534 ddname OPEN FAILURE--AVAILABLE COMMANDS: HELP,END**
Reason: The required library on DD statement "ddname" could not be opened. Commands which access modules in the that library cannot be executed.
Action: Check the joblog for IBM OPEN error messages. Correct the error and rerun.
- FDR535 mmmmmmmm READ UNSUCCESSFUL--function BYPASSED**
Reason: Program attempted to read a member from a library or a data set but the read was not successful. "mmmmmmmm" can be:
module MODULE - if the error occurred trying to read a load module
MEMBER - if the error occurred trying to read another kind of member
DATA SET - if the error occurred trying to read a sequential data set or member.
Action: See message FDR530.
- FDR536 member CONTAINS NO EMPTY SLOTS--UPDATE BYPASSED**
Reason: An attempt to add more definitions to an ABR protect list failed because there was no room for new entries in that list.
Action: Display the protect list module with ISPF panel A.1.6 and eliminate duplicate or obsolete entries. If this is not possible, the FDR Installation Control Library (ICL) contains the source code for these lists, allowing you to expand them for more entries.
- FDR537 member CONTAINS NO PROTECTED DSNAMES/DSGROUPS/VOLUMES**
Reason: FDRZAPOP attempted to print entries from the ABR protect list module "member" but the member is empty.
- FDR538 MODULE member WRONG VER/LEVEL--ddname=dsname**
Reason: Program read module "member" in the library "dsname" specified on the DD statement "ddname", but the module loaded is not at the same level as the calling program.
Action: Verify that the data set name specified on the DD statement was the correct one. If incorrect, change and resubmit the job. If the data set name is correct, check to see if all the installation steps required were done ([see Section 90](#)). Contact Innovation for technical assistance.

100.03 CONTINUED . . .

- FDR539** **MODULE *member* function--*ddname*=*dsname***
Reason: The module "member" in the library "dsname" specified on the DD statement "ddname" has been successfully processed. "function" indicates what operation was done.
- FDR540** ***operand* CONTAINS INVALID CHARACTERS--
ZAP REJECTED**
Reason: The value specified for "operand" contained one or more characters that were not A-Z, 0-9, \$#@.
Action: Correct the value and resubmit the job.
- FDR541** **ALLINDX CREATES *name* NAME GT *nn* CHARACTERS--
ZAP REJECTED**
Reason: The value specified for the operand ALLINDX, when used to replace the first index of the option "name", creates a character string longer than the limit "nn".
Action: Shorten the value of ALLINDX by specifying fewer characters and resubmit the job. Some options may have already been changed before the message is issued.
- FDR542** ***operand* CONTAINS INVALID INDEX STRUCTURE--
ZAP REJECTED**
Reason: The value specified for "operand" contains two (2) or more consecutive periods (..) in violation of MVS data set naming conventions.
Action: Correct the error and resubmit the job.
- FDR543** ***operand* DOES NOT CONTAIN CHARACTER STRING *rrrrrr*--
ZAP REJECTED**
Reason: The value specified for "operand" does not contain the character string "rrrrrr" as an index level. If "rrrrrr" is ARCHIVE, then valid ARCDSDN values would be 'SAM.ARCHIVE' or 'FDRABR.CURRENT.ARCHIVE.DATASET'.
Action: Provide a proper data set name and resubmit the job.
- FDR544** ***operand* REJECTED--EXCEEDS MODIFIABLE PORTION OF FDROPT**
Reason: The length of the value specified for "operand" taken in conjunction with the offset operand exceeds the modifiable portion of FDROPT.
Action: Contact Innovation for assistance.
- FDR545** **VERIFY FAILED--CHAR/HEX PRINT FORCED**
Reason: The VERIFY of existing contents failed. A character/hex print of the module FDROPT is produced. Always preceded by message FDR544.
- FDR546** **CYCLE TABLE FIELD EXCEEDED -- MIN OFFSET -- *nnn*/MAX ENTRIES -- *eee***
Reason: The length of the value specified for the operand OLDBENT, taken in conjunction with the value of the operand OLDBOFF exceeds the limits of the OLDBACKUP table.
Action: Correct the error(s) and resubmit the job.
- FDR547** **FDRCPKUM UNMOVABLE TABLE HAS ONLY *nnn* REMAINING BYTES -- UPDATED BYPASSED**
Reason: An attempt to add more definitions to the COMPAKTOR unmovable data set table failed because there was no room for the new entries in that table. "nnn" bytes remain unused; it might still be possible to add entries that are smaller than the failing entry.
Action: Display the unmovable table module with ISPF panel A.1.6 and eliminate duplicate or obsolete entries. If this is not possible, the FDR Installation Control Library (ICL) contains the source code for this table, allowing you to expand it for more entries.
- FDR548** **FDRCPKUM CONTAINS NO UNMOVABLE DATA SET NAMES/GROUPS**
Reason: FDRZAPOP attempted to print entries from the COMPAKTOR unmovable table but the table is empty.
- FDR549** **PRINTING *module*--SYSLIB=*dsname***
Reason: The module "module" from the data set referenced by ddname SYSLIB is being printed as requested by the user.
- FDR550** **ERRORS ENCOUNTERED DURING EXECUTION--REWRITE CANCELLED**
Reason: One or more modules from the SYSLIB data set have been updated and were scheduled to be rewritten to the SYSLIB data set. However, previous commands failed to complete successfully so the updates were discarded and the modules were not modified.
Action: Correct the error conditions documented by the error messages and resubmit the job.

100.03 CONTINUED . . .

- FDR552** **MAXIMUM VOLSER COUNT OF *nnnnn* EXCEEDED--DUPLICATE VOLSER MAY BE PRINTED**
Reason: During 'PRINT VOLSTAT' processing in FDRABRP the maximum number of disk volumes was exceeded. In rare cases this might result in a volume being printed more than once. Additional volumes are not processed.
Action: Increase the value of MAXONLINE in the FDR Global Option Table, using ISPF option A.I.4.10. The default is 256.
- FDR553** **VOLSER TABLE FULL -- ENTRY COUNT IS *nnnnn* -- NO ADDITIONAL ENTRIES WILL BE DROPPED**
Reason: A REORG of the ARCHIVE control file has been requested but the maximum number of Archive Backup tapes to be tabled has been exceeded. When no more tapes can be tabled, no more entries will be dropped from the ARCHIVE control file.
Action: Rerun the REORG as soon as possible, specifying the operand MAXTREORG=*nnnnn* (100 to 32000) with a value greater than the number of the entries shown in the message. The value of MAXTREORG can be changed permanently in the FDR/ABR global option table on ABR ISPF panel A.I.4.6.
- FDR554** **NULL *statement* SUBCOMMAND--SUBCOMMAND IGNORED**
Reason: "*statement*" was input with no operands. One or more operands are required.
Action: Correct the statement and resubmit the job.
- FDR555** ***operand*/OFFSET MUST BOTH BE SPECIFIED--SUBCOMMAND IGNORED**
Reason: "*operand*" must be used in conjunction with OFFSET.
Action: Specify the missing operand and resubmit the job.
- FDR556** ***operand*/OFFSET OF *nnn* EXCEEDS LENGTH OF *record* -- SUBCOMMAND IGNORED**
Reason: The combined length, either implied or specific, of "*operand*" and the OFFSET operand exceeded the length of the data to be modified, identified by "*record*".
Action: Correct the error and resubmit the job.
- FDR557** ***operand*/OFFSET DO NOT CORRESPOND--SUBCOMMAND IGNORED**
Reason: The user specified an unequal number of "*operand*" and OFFSET operands.
Action: Correct the error and resubmit the job.
- FDR558** **FIELD NAME *fieldname* NOT ELIGIBLE AS SORT CONTROL -- SUBCOMMAND IGNORED**
Reason: "*fieldname*" cannot be used as a sort field.
Action: Remove the offending field name from the sort field selection list and resubmit the job.
- FDR559** **FIELD NAME *fieldname* *reason***
Reason: "*fieldname*" was not properly specified for "*reason*". In most cases, the field name is ignored but the program continues.
Action: Remove the offending field name from the report field selection list and resubmit the job.
- FDR560** **INPUT ARCHIVE OPEN/VALIDATION ERRORS--MERGE SUBCOMMAND CANCELLED**
Reason: While opening input files for an Archive Control File MERGE, FDRARCH discovered errors either opening or validating the names or contents of one or more of the input ARCHIVE files. Other messages will define the exact errors.
Action: See the action for the other error messages.
- FDR561** **FIELD NAME *fieldname* REJECTED FOR SORT/BREAK - MULTIPLE RECORD FORMATS NOT PERMITTED.**
Reason: "*fieldname*" was rejected on a SORT or BREAK statement because it comes from a different source than other field names already specified.
Action: Correct the error and resubmit the job.
- FDR562** **FIELD NAME *ffffff* IS AN UNDEFINED USER FIELD -- SUBCOMMAND IGNORED**
Reason: FDREPORT has nine (9) pre-defined user fields that must be activated and redefined prior to use. This field has not been activated.
Action: display the HELP for the ACTIVATE command using this control statement: HELP CO(ACTIVATE)
- FDR567** **OBSOLETE OPERAND *operand* SPECIFIED--CONVERTED TO *newoperand***
Reason: User specified "*operand*" which is obsolete. Rather than bypass the command, the operand was converted internally to a currently supported operand "*newoperand*" that will perform the same function.
Action: For future use, you could convert the job to use the new operand.

100.03 CONTINUED . . .

- FDR568 REQUIRED OPERAND *operand* MISSING-- status**
Reason: "operand" is required but was not provided (it may list more than one operand). "status" lists the action taken. If "status" ends in IGNORED, then it identifies another operand which was ignored. If it ends in ASSUMED, it displays assumed defaults.
Action: If necessary, correct the statement and resubmit the job.
- FDR569 WARNING -- option NOT CATALOGED--DSN=*dsname***
Reason: The FDR Global Option Table was updated with a new data set name, for example, the Archive Control File name or the Remote Queue data set names. The new data set name was not in the system catalogs and may not exist.
Action: This is only a warning message. The data set name listed has become the new name of the indicated default.
- FDR574 reason--DDNAME=*ddname*--status**
Reason: "ddname" could not be opened for "reason", resulting in "status". If the DD is required, the function will be terminated. In some cases, another action is taken and the function will continue.
Action: Check the joblog for IBM OPEN error messages or misspelled DD names. Correct and resubmit the job.
- FDR575 FAMS ACCESS ERROR -- description -- DSN=*dsname***
Reason: An access request was made via the FAMS facility to gather information about a PDSE named "dsname". The access request failed; "description" may describe the error in more detail.
Action: Contact Innovation for assistance.
- FDR577 IAM LOCATE ERROR -- description -- DSN=*dsname***
Reason: A special form of the LOCATE SVC was issued through the IAM/VIF interface to gather information about an IAM data set named "dsname". The locate failed with an unexpected return code. "description" may describe the error in more detail.
Action: Contact Innovation for assistance.
- FDR580 LSPACE ERROR -- VOLSER=*volser* -- RETURN CODE=*rr* DADSM DIAGNOSTIC INFO=*dddddddd***
Reason: An LSPACE SVC was issued to get free and allocated space information for volume "volser" but it failed with return code "rr" in decimal and diagnostic information "dddddddd" in hex.
Action: Contact Innovation for assistance.
- FDR585 RECORD SIZE *nnnnn* EXCEEDS MAXIMUM DEVICE BLOCK SIZE**
Reason: The number of data bytes necessary to construct the desired output record from FDREPORT exceeds the maximum device block size on the device assigned to the report data set.
Action: Change to a different output device or remove some of the field names to reduce the size of the output record.
- FDR586 service CANCELLED -- AUDIT CONTINUES**
Reason: User has selected an FDREPORT output service that is not available, usually because of a missing or incorrectly specified DD statement. This message will probably be preceded by an FDR465 detailing the DDname in error.
Action: See message FDR465.
- FDR587 SUMMARY VOLSER TABLE FULL -- ENTRY COUNT IS 5000 -- NEW VOLUMES WILL NOT BE SELECTED**
Reason: DISABLE=NEWSUMFORMAT has been specified by the user along with SUM=YES and DATATYPE=EXTRACT or DATATYPE=ARCHIVE. The selection criteria specified by the user caused more than 5000 disk volumes to be tabled for summary processing.
Action: Processing of this volume is bypassed. Volumes tabled prior to receiving this error message will be processed. For a more complete output, rerun the job specifying ENABLE=NEWSUMFORMAT. Old summary processing is limited to 5,000 disk volumes when processing DATATYPE=EXTRACT or DATATYPE=ARCHIVE.
- FDR588 VOLSER SELECTION TABLE FULL -- MAXONLINE VALUE IS *nnnn* -- VOLSER *volser* BYPASSED**
Reason: The selection criteria specified by the user caused more than "nnnn" disk volumes to be tabled for processing. For example, DATATYPE=CATARCH or DATATYPE=CATVTOC attempted to table all volume serial number of each and every cataloged data set selected.
Action: Processing of volume "volser" is bypassed. Volumes tabled prior to receiving this error message will be processed. For a more complete output, rerun the job specifying the MAXONLINE operand with a value that is greater than the value printed.

100.03 CONTINUED . . .

- FDR589 VOLSER SELECTION TABLE FULL --
MAXONLINE VALUE IS nnnn -- REQUEST CANCELLED**
Reason: The selection criteria specified by the user caused more than "nnnn" disk volumes to be tabled for processing. For example, UNITNAME=SYSALLDA caused the volume serial number of each and every mounted disk volume to be tabled or more than "nnnn" disk volumes were specified by the VOL= operand.
Action: Processing is terminated prior to performing the specified request. Rerun the job, specifying the MAXONLINE operand with a value that is greater than to the number of online disk volumes in the installation.
- FDR596 STORAGE VALUE OF nnnnnnn INSUFFICIENT FOR MERGE --
SUBCOMMAND TERMINATED**
Reason: The MERGE work area is not large enough to hold all the date ranges required to MERGE the Archive Control Files.
Action: Resubmit the job specifying a value for the STORAGE operand greater than the value printed. It is suggested the value be at least doubled.
- FDR598 ALLFILES/MAXFILES=/EXP= NOT ALLOWED
WITH DISK OR NON-SL TAPE INPUT**
Reason: The COPY statement input to FDRTCOPY specifies either the ALLFILES, MAXFILES=, or EXP= operands, but the TAPEIN DD statement does not point to a tape, or is a unlabeled tape.
Action: Correct the COPY statement or the DD statement.
- FDR603 status BACKUP OF DSN=dsname**
Reason: This message is displayed by FDRTCOPY if the PRINT=DSN operand is coded, or by FDRtsel with ARCEDIT. It indicates that it is **COPYING**, **DROPPING**, or **KEEPING** the backup of data set "dsname"..
- FDR604 ERROR occurred POSITIONING TAPEIN REASON CODE-x**
Reason: FDRTCOPY detected an error when positioning the input tape during ALLFILES processing. The reason code "x" may be:
1 -- Tape mark not read when expected.
2 -- Tape mark read when not expected.
3 -- Trailer label read is not EOVS1/EOFS1.
4 -- Header label read is not HDR1.
5 -- Tape drive unit check during positioning.
6 -- Other I/O error during positioning.
Action: A U0659 ABEND follows. This may be due to hardware problems with the tape drive. If it occurs with the same input tape on more than one tape drive, contact Innovation for assistance.
- FDR605 COPY TERMINATED AT TAPE MARK SEQUENCE**
Reason: During a COPY with ALLFILES or MAXFILES=, FDRTCOPY found two consecutive tape marks following the trailer labels of the last file processed, indicating the end of the input tape(s).
Action: None. This is a normal completion.
- FDR606 COPY TERMINATED BY FILE COUNT**
Reason: During a multi-file COPY, FDRTCOPY completed copying the number of files specified by the MAXFILES= operand.
Action: None. This is a normal completion.
- FDR607 COPY TERMINATED BY END OF VOLUME LABELS**
Reason: During a COPY with ALLFILES or MAXFILES=, FDRTCOPY detected that the last file copied was followed by EOVS (End-Of-Volume) labels rather than EOF (End-Of-File) labels, so FDRTCOPY is not able to continue. This probably means that the last file copied is not complete, which will result in a FDR203 message
Action: None. This is not an error condition unless an FDR203 message is also printed to indicate that the last file copied was not a complete FDR backup.

100.03 CONTINUED . . .

- FDR609 CURRENT PAGEWIDTH(*nnn*) EXCEEDED BY *vvv*--action**
Reason: An FDREPORT step specified PAGEWIDTH=*nnn* but the total length of the fields requested to be printed exceeds this value. If the output report DDNAME is not SYSPRINT, the output DCB is opened with a LRECL that corresponds to the pagewidth. The result is truncation of the print line and the action will be 'PRINT LINE TRUNCATED'. If the report DDNAME is SYSPRINT, 'PRINT LINE TRUNCATED' will appear only if the actual LRECL of 121 is exceeded.
Action: Increase the PAGEWIDTH, specify the AUTOSTACK operand to force eligible field stacking, or reduce the number of fields to be printed.
- FDR610 ICF CLUSTER TABLE FULL--ENTRY COUNT IS *nnnnn*--ICF SELECTION IS TERMINATED**
Reason: Selection of ICF VSAM data sets is being performed by base cluster name utilizing a look-aside table of clusters. The table has been filled, and selection of additional ICF VSAM data sets is no longer performed.
Action: See message FDR670.
- FDR615 dddddddd DSN=*dsname* FILE=*fn* VOL=*vvvvvv* status**
Reason: Printed by FDRTCOPY at the end of each file copied. One FDR615 message will be printed for TAPEIN, TAPEOUT, and TAPE2OUT (if present). It documents the input or output DDNAME, data set name, file number (0 if on disk), and up to 20 volume serials. "status" will be **DROP** if an input file was not copied due to being expired, or **CATLG** if an output file was successfully cataloged.
- FDR616 FDRARCH ABNORMALLY TERMINATED. COMP CODE=*Ssss Uuuuu***
Reason: FDRTCOPY invoked FDRARCH to update the Archive Control File with information about copied Archive Backups, but FDRARCH abnormally terminated. The termination code is either system ABEND "sss" or user ABEND/return "uuuu".
Action: Examine FDRARCH messages on SYSPRINT to determine the cause. The Archive Control File update may have been partially or totally bypassed. If possible, correct the problem and re-execute the FDRTCOPY job. If not, contact Innovation for technical assistance. U0012 indicates that FDRARCH terminated with return code 12 due to some error; look for preceding FDRARCH error messages.
- FDR618 PROCESSING VOLUME *vvvvvv* AS status**
Reason: FDREPORT tried to bring the used portion of a VVDS into storage to satisfy ICF/VSAM report requests, but was unable to read the VVDS. An FDR63x message documenting the cause will precede this message.
Action: If "status" is **ICF=IGNORE**, the VVDS is unusable and ICF VSAM components will be processed as if they were non-VSAM data sets. This may indicate a serious problem with this volume.
 If "status" is **ICFSOURCE=LOCATE**, catalog LOCATEs will be used to get information on ICF VSAM components and clusters instead of reading the VVDS directly. If no other errors occur, FDREPORT will run successfully, but it will run longer.
- FDR619 MODULE *c...c* TOO LARGE TO PROCESS-SYSLIB=*d...d***
Reason: FDRZAPOP attempted to read the module *c...c* in the data set *d...* referenced by SYSLIB. The module was required to support all processing functions but, due to the size of the module (or previously read modules), not enough buffer storage was available to complete the read. Because of user expansion of the protect, allocation, and/or unmovable tables, the storage buffer used by FDRZAPOP to store tables has been exceeded.
Action: Contact Innovation for assistance in expanding the storage buffer used by FDRZAPOP.
- FDR620 statement error**
Reason: In FDRTSEL, "statement" was in error. "error" can be **COMMAND MISSING OR INVALID** or **COMMAND PREVIOUSLY SPECIFIED**.
Action: Correct the error and reexecute.
- FDR621 SELECT COMMAND FUNCTION MISSING. MUST SPECIFY ARCHIVE, ARCEDIT, OR CATLG**
Reason: In FDRTSEL, the SELECT statement does not contain one of the required operands ARCHIVE, ARCEDIT or CATLG.
Action: Correct the error and reexecute.

100.03 CONTINUED . . .

- FDR622 SELECT COMMAND COPY= OR IFONLYCOPY= KEYWORD IS MISSING OR MISPELLED**
Reason: In FDRSEL, the SELECT statement does not contain one of the required COPY= or IFONLYCOPY= operands.
Action: Correct the error and reexecute.
- FDR623 REASON=reason**
Reason: This message documents various errors which prevent FDRSEL from completing. "reason" may be:
A - operand IS INVALID WHEN SELECT COMMAND IS PRESENT
 The indicated operand of the COPY or MOVE statement is invalid with a SELECT statement.
B - DDNAME=TAPE2OUT IS INVALID WITH TAPEOUT TO DISK
 If TAPEOUT points to a disk, TAPE2OUT cannot be present.
C - DDNAME=TAPEOUT# IS INVALID IN JCL FOR FDRSEL
 TAPEOUT# is a reserved Ddname.
E - MOVE COMMAND INVALID WITHOUT A TAPEOUT DD - COPY FORCED
F - MOVE COMMAND CAN NOT BE SPECIFIED WITHOUT A SELECT COMMAND
G - MOVE COMMAND INVALID WITH TAPE INPUT. CHANGE TO COPY OR MODIFY BKDEVTYPE
 The MOVE statement can not be specified when the TAPEIN file is not on DISK. When the selected files are on tape, only COPY is valid.
H - SYSIN ERROR - READING PAST EOF
 Should not occur. Contact Innovation Technical Support if you receive this error.
J - CAT=NO SPECIFIED WITH DISK OUTPUT NOT ALLOWED
 CAT=NO must not be specified when the TAPEOUT files are being directed to DISK on a disk-to-disk copy.
Action: Correct the error and reexecute. Consult [Section 60](#) for FDRSEL documentation.
- FDR624 INVALID KEYWORD VALUE POINTED BY ? (QUESTION MARK) ON THE ABOVE STATEMENT**
Reason: FDRSEL detected an invalid keyword value. The preceding FDR624 message contains a ? (question mark) under the invalid keyword value.
Action: [See Section 60](#) for information on the supported keyword values.
- FDR625 NO FILES SELECTED. NO FILES MATCH THE SPECIFIED SELECT CRITERIA**
Reason: The specified selection criteria does not match any recorded ABR backup.
Action: This may be a normal result; if not, run the PRINT ARCHIVE documented in [Section 53](#) to inspect the records in the archive file.
- FDR626 ARCBACKUP=DSF DUMPING FILE: dsname**
Reason: FDRSEL found the ARCBACKUP=DSF operand was specified so it will invoke FDRDSF to backup the updated Archive Control File on the end of the current tape.
- FDR627 DSN=dsname function status**
Reason: FDRSEL performed the indicated function on the Archive Backup "dsname". "function" may be **UNCATALOGED**, **CATALOGED**, **SCRATCHED** or **RENAMED**. "status" will be blank or **FAILED** if the function failed.
- FDR628 LOCATE FOR GDG FAILED FOR DSN=dsname**
Reason: FDRSEL received a non-zero return code from a catalog LOCATE for the Archive Backup "dsname". This message is accompanied by a FDR516 message containing the error return codes.
Action: See message FDR516.
UNCATALOG FAILED FOR DSN=dsname
Reason: In this form of message FDR628, FDRSEL received a non-zero return code for the request to uncatalog an Archive Backup. This message is accompanied by an FDR516 message containing the error return codes, and by another FDR629 message containing the scratch return code for each of the volumes where the archive backup resides.
Action: See message FDR516.
- FDR629 DSN=dsname SCRATCH FAILED volser CODE=nnnn**
Reason: FDRSEL received a non-zero return code for the request to scratch a disk Archive Backup. This message is accompanied by an FDR516 message containing the error return codes, and by another FDR629 message containing the scratch return code for each of the volumes where the archive backup resides.
Action: See message FDR516.

100.03 CONTINUED . . .

- FDR630 INSUFFICIENT STORAGE *service description* -- AMOUNT REQUIRED nnnnnnnn BYTES**
Reason: FDREPORT requires additional storage to perform the service described but is unable to get enough storage.
Action: and resubmit the job. If message FDR631 follows, FDREPORT will run successfully.
- FDR631 VTOC WILL BE READ ONE (1) TRACK AT A TIME**
Reason: Because of insufficient storage, FDREPORT was unable to read the entire VTOC into storage, so it will be read one track at a time.
Action: FDREPORT should complete successfully, but the elapsed time will be increased. Increase the region size for future use.
- FDR632 SYS1.VVDS.Vvvvvvv IN ERROR -- CONTAINS EXTENT COUNT OF ZERO
 SYS1.VVDS.Vvvvvvv REQUIRED -- NAME NOT FOUND IN VTOC**
Reason: FDREPORT was trying to read the VVDS since the volume contains ICF VSAM clusters or SMS-managed data sets, but it failed for the reason indicated.
Action: message FDR618 will follow.
- FDR633 SYS1.VVDS.Vvvvvvv IN ERROR -- FORMAT 3 DSCB NOT FOUND IN VTOC**
Reason: FDREPORT was trying to read the VVDS since the volume contains ICF VSAM clusters or SMS-managed data sets, but it failed for the reason indicated.
Action: message FDR618 will follow.
- FDR634 SYS1.VVDS.Vvvvvvv IN ERROR -- I/O ERROR READING VVDS**
Reason: FDREPORT was trying to read the VVDS since the volume contains ICF VSAM clusters or SMS-managed data sets, but it failed for the reason indicated.
Action: message FDR618 will follow.
- FDR635 TRKCALC FAILED -- R15=xxxxxxx -- VOL=vvvvvv -- DSN=dsname**
Reason: The TRKCALC SVC was issued for the data set name listed to determine the number of records that will fit on a track of the device. The SVC failed with a non-zero return code in register 15.
Action: The following fields will be set to zero (0) as a result of the failing TRKCALC:
 %CAPUSED, BLKSTRK, BYTESFRE, BYTESUSE, CAPBYTES
 Call Innovation for support.
- FDR636 UNABLE TO DETERMINE VERSION NUMBER/LEVEL -- LRECL=nnnnn**
Reason: An FDREPORT extract file was input to FDREPORT (DATATYPE=EXTRACT). The extract file can have several formats depending on options used and the version of FDREPORT used to created it, but FDREPORT was unable to identify the format. The logical record length of the failing file is "nnnnn".
Action: If the extract file was created by a higher version of FDREPORT than the version trying to read it, use the higher version. Otherwise, contact Innovation for assistance.
- FDR637 CONVERTING EXTRAC FILE FROM VERSION NUMBER/LEVEL xxxxxxxx**
Reason: An FDREPORT extract file was input to FDREPORT (DATATYPE=EXTRACT). The extract file can have several formats depending on options used and the version of FDREPORT used to created it. FDREPORT has identified that the format is that created by an earlier version of FDREPORT.
- FDR639 PRINT REQUEST BYPASSED -- DATATYPE REQUIRES APF AUTHORIZATION -
 FDREPORT NOT APF AUTHORIZED.**
Reason: FDREPORT requires APF authorization to complete the report you requested, but it is not authorized.
Action: Place FDREPORT in a APF-authorized library and resubmit the job.
- FDR640 UNIT NAME uuuuuuuu *reason***
Reason: A request was made to report on volumes selected by logical unit name (UNIT=uuuuuuuu) but the attempt to select volumes failed for "reason".
Action: The request is terminated.
- FDR641 STORAGE GROUP=storgroup *reason***
Reason: A request was made to report on volumes selected by SMS storage group name (STORGRP=storgroup) but the attempt to select volumes failed for "reason".
Action: The request is terminated.

100.03 CONTINUED . . .

- FDR642 service NOT SUPPORTED - FDREPORT NOT APF AUTHORIZED.**
Reason: FDREPORT requires APF authorization to honor the service you requested, but it is not authorized.
Action: Place FDREPORT in a APF-authorized library and resubmit the job.
- FDR643 command SUBCOMMAND IGNORED**
Reason: This is always preceded by another message indicating why the command could not be processed.
- FDR648 I/O ERROR location -- VOL=vvvvvv -- DSN=dsname**
Reason: An I/O error occurred during EXCP processing to extract information from either a PDS directory or a suspected IAM data set.
Action: This message is normally the only documentation produced when these error conditions occur. Further documentation in the form of the standard Innovation mini-dump will be produced if ENABLE=DEBUG is specified on either the DEFAULT or PRINT commands. If the error persists, obtain all available documentation and call Innovation .
- FDR650 ASTERISK TRUNCATION NO LONGER VALID FOR UNIT**
Reason: Earlier versions of FDREPORT supported an asterisk at the end of the value of the UNIT parameter, e.g., UNIT=13*. Because of the support for 4-digit device addresses, this is no longer supported.
Action: If the comparison operand was an equal sign (e.g., UNIT=13*), it will be converted into a range of unit addresses; see message FDR652.
- FDR651 UNIT unit CONTAINS INVALID CHARACTERS ONLY 0-9 AND A-F ARE VALID**
Reason: A UNIT= parameter had a value which was not a valid hexadecimal device number
Action: Correct and resubmit
- FDR652 OBSOLETE UNIT=unit -- CONVERTED TO UNIT>=uuu0,UNIT<=uuuF**
Reason: Earlier versions of FDREPORT supported an asterisk at the end of the value of the UNIT parameter, e.g., UNIT=13*. Because of the support for 4-digit device addresses, this is no longer supported. But because the comparison operator was an equal sign (UNIT=) FDREPORT was able to convert to the two UNIT values shown in the message.
Action: You should update your FDREPORT input to use similar UNIT parameters
- FDR653 SYSUT2 ALLOCATION ERROR -- EXTRACTMEMBER REQUIRES PARTITIONED DATA SET (DSORG=PO)**
Reason: EXTRACTMEMBER=membername was specified on a FDREPORT PRINT statement, but the data set identified by EXTRACTDSNAME=dsname was not a partitioned data set.
Action: If the input data set is really sequential, omit EXTRACTMEMBER=. If EXTRACTDSNAME= identifies the wrong data set, correct it.
- FDR654 EXTRACTDSNAME SPECIFICATION ERROR -- description**
Reason: EXTRACTDSNAME=dsname was specified on a FDREPORT PRINT statement, but "dsname" was not a proper data set name, for the "reason" displayed.
Action: Correct and resubmit.
- FDR655 NUMBER OF TAPExxxx DD STATEMENTS EXCEEDS MAXIMUM**
Reason: For an FDREPORT PRINT with DATATYPE=TVTOC, the number of TAPExxxx DD statements specifying FDR backups to read exceeded 256.
Action: If possible, split the step into multiple steps. If not, contact Innovation for assistance in raising the limit.
- FDR656 DDNAME=TAPExxxx IS AN ABSOLUTE TRACK DUMP DATA SET -- UNABLE TO PROCESS**
Reason: For an FDREPORT PRINT with DATATYPE=TVTOC, the TAPExxxx DD shown contained a dump created with the absolute track backup feature of FDRDSF. There is no VTOC information on it to display, so it is bypassed.
- FDR657 PREMATURE END OF FILE ON DDNAME=TAPExxxx -- PROCESSING SUSPENDED**
Reason: For an FDREPORT PRINT with DATATYPE=TVTOC, the TAPExxxx DD shown contained an FDR backup, but the end of the backup was encountered before all of the FDR control records were read. Data from control records which were read will be included in the report.
Action: If the backup is on multiple tape volumes, be sure that all those volumes were specified on the TAPExxxx DD, and in the correct order.

- FDR658 LSTAR ERROR -- USED SET TO ALLOCATED -- VOL=volser DSN=dsname**
- Reason:** For the identified data set, FDREPORT found that the DS1LSTAR (last block pointer) in its Format 1 DSCB was invalid. The number of used tracks (USED) is set to the number of allocated tracks (SIZE).
- FDR660 RECLAIMED FROM DELETION -- DSNAME=dsname**
Reason: For the identified data set, FDRARCH had previously indicated that the GDG generation named was deleted from the Archive Control File during a reorganization. However, to meet the requirements of MAXGENERATIONS=nnn and ENABLE=RECLAIMGEN, it was kept instead.
- FDR661 RETAINED -- CATERROR=RETAIN SPECIFIED -- DSNAME=dsname**
Reason: A catalog error occurred when trying to delete the auto-recall catalog entry for "dsname". The data set was not deleted from the Archive Control File.
- FDR662 NOT CATALOGED ASSUMED -- CATERROR=option SPECIFIED**
Reason: A catalog error occurred when trying to delete the auto-recall catalog entry for "dsname". Either CATERROR=IGNORE or CATERROR=PURGE was specified, so the data set is treated as if it were not cataloged
Action: You may need to manually delete the auto-recall catalog entry.
- FDR670 INSUFFICIENT STORAGE FOR ICF CLUSTER PROCESSING -- AMOUNT REQUIRED nnnnnnnn BYTES**
Reason: A program attempted to GETMAIN nnnnnnnn bytes of additional storage to table the active clusters within the Archive Control File but the storage was not available.
Action: Increase region size and resubmit the job. Always followed by message FDR610 indicating ICF cluster processing is terminated for the remainder of this command execution.
- FDR671 ddname REFERENCES AN EMPTY ARCHIVE CONTROL FILE -- subcommand CANCELLED**
Reason: The ARCHIVE Control File being processed contains no records (i.e.: it has only been formatted by FDRARCH and never used by FDRABR).
Action: Verify that the proper data set name was used for the ARCHIVE Control File you wished to process and correct if required.
- FDR672 ARCHIVED ON yyyy.ddd - FROM volser - DROPPED - DSN=dsname**
Reason: The entry for data set name "dsname" archived from disk "volser" on Julian date yy.ddd was dropped from the ARCHIVE Control File. This message will always be followed by one (1) or more FDR673 messages detailing the reason the entry was dropped and other action taken on behalf of the named data set or component.
- FDR673 action or reason description**
Reason: The following are the descriptions that can be displayed:
COMPONENT OF DROPPED CLUSTER -- CLUSTER=cluster name
COMPONENT HAS NO RELATED CLUSTER
ENTRY IS A DUPLICATE DATA SET NAME -- MAXOCCURANCES=nnnnn
ENTRY IS AN EXCESSIVE GENERATION -- MAXGENERATIONS=nnnnn
ENTRY MARKED AS HAVING BEEN RESTORED -- RESTORE=PURGE SPECIFIED
ENTRY MARKED AS HAVING BEEN DELETED -- DELETE=PURGE SPECIFIED
ALL ARCHIVE DUMP TAPES HAVE EXPIRED -- EXPIRE=PURGE SPECIFIED
ENTRY NOT MARKED FOR AUTO RECALL -- ENABLE=IFNOTAUTOREC SPECIFIED
ENTRY NOT CATALOGED FOR AUTO RECALL -- ENABLE=IFNOTCAT SPECIFIED
DATA SET ON DISK -- AUTO RECALL INDICATOR REMOVAL FAILED
DATA SET ON DISK -- AUTO RECALL INDICATOR REMOVED
UNCATALOGED FROM AUTO RECALL
- FDR674 RETURN CODE rrrr REASON IGG0CLcc-eeee**
Reason: A DELETE NOSCRATCH failed for the data set name listed in the FDR511 or FDR514 message that immediately precedes this message. The return code "rrrr" and reason code are the same as documented for message IDC3009I in the IBM System Messages manual. IGG0CLcc is the name of the IBM catalog module which generated the error.
- FDR675 CONTROL FILE CONTAINS nnnnnnnn ACTIVE RECORDS**
Reason: List the current record count of the Archive Control File identified by messages FDR505 and FDR507 when these messages are issued by the DUMP or REORG subcommands.

100.03 CONTINUED . . .

- FDR676** **nnnnnn VOLUME(S) OF DEVICE=*devicetype* WILL BE RELEASED**
Reason: The reorganization of the Archive Control File has resulted in "nnnnnn" tape volumes of the device type listed being made available for scratch use. They were required for the restore of Archived data sets that have been dropped from the Archive Control File. If this was a simulation, this documents the volumes that will be released.
- FDR677** **CAUTION -- *function* REQUIRES EXECUTION OF THE RECATALOG SUBCOMMAND WITH SERVICE=MERGE TO CORRECT THE AUTO RECALL CATALOG ENTRIES**
Reason: This 3 line message calls attention the fact that you may need to execute FDRARCH RECATALOG SERVICE=MERGE to correct auto-recall catalog entries. "function" can be:
BLOCK SIZE CHANGE RESULTING FROM THE RESTORE SUBCOMMAND
PROCESSING PERFORMED BY THE MERGE SUBCOMMAND
Action: If this is to be the new permanent Archive Control File, re-execute FDRARCH against this control file with the following subcommand:
 RECATALOG SERVICE=MERGE
 If this is only a test this message may be ignored.
- FDR678** **SORT=NO SPECIFIED -- THE FOLLOWING OPERANDS WILL BE IGNORED**
Reason: SORT=NO was specified on the REORG subcommand of FDRARCH along with operands that are only valid if SORT=YES is specified. The list of the operands that will have no effect on this reorganization are listed on subsequent lines.
Action: If the operands that were ignored are required for your installation, re-execute FDRARCH changing SORT=NO to SORT=YES. SORT=YES is recommended for REORG.
- FDR679** **SORT=YES FORCED -- CAUSED BY THE FOLLOWING OPERANDS**
Reason: The SORT= operand was not specified on the REORG subcommand of FDRARCH, although one or more operands were specified that are only valid if SORT=YES is specified. It is assumed that SORT=YES was IMPLIED by specifying the operands that are listed on subsequent lines.
- FDR680** **CATALOG SCAN USING FDR\$CATP NOT PROCESSED - CATALOG ACCESS REQUIRES APF AUTHORIZATION.**
Reason: FDREPORT requires APF authorization to complete the report you requested, but it is not authorized.
Action: Place FDREPORT in a APF-authorized library and resubmit the job.
- FDR681** **REORG SUBCOMMAND USING SORT NOT PROCESSED -- *reason***
Reason: The SORT= operand was specified on the REORG subcommand of FDRARCH but the SORT-based REORG could not be executed due to errors involving multi-level alias (MLA) in the catalog. "reason" may be:
MULTI-LEVEL ALIAS REQUIRES FDRARCH BE APF AUTHORIZED
CATALOG ADDRESS SPACE NOT FOUND
CATALOG ADDRESS SPACE NOT FOUND
Action: If FDREPORT is not executing as an authorized program, make it so. Contact Innovation for support
- FDR682** **CANCEL PROTECT TASK STARTED -- *hh.mm.ss***
Reason: Some functions of FDRARCH must be allowed to complete to avoid damage to the Archive Control File, so CANCEL protection is used for them. CANCEL protection intercepts operator CANCEL commands and asks the operator to confirm or ignore the CANCEL. The CANCEL protected task was started at the time indicated.
- FDR683** **CANCEL PROTECT TASK ENDED -- *hh.mm.ss***
Reason: See FDR682 The CANCEL protected task was ended at the time indicated.
- FDR684** ***description* SERVICE POSTED COMPLETE -- *hh.mm.ss***
Reason: See FDR682 The CANCEL protected task has completed at the time indicated.
- FDR685** **CANCEL PROTECT TASK COMPLETION CODE - *nnnnnnnn***
Reason: See FDR682 The CANCEL protected task completed with the return code "nnnnnnnn".
- FDR686** **SERVICE TERMINATED IN RESPONSE TO A STOP REQUEST**
Reason: See FDR682 The operator replied "STOP" to the CANCEL protection console message, so FDRARCH will bring the current function to a graceful stop.
- FDR687** **CATALOG SCAN USING FDR\$CATP NOT PROCESSED - UNABLE TO CREATE ALIAS AND CATALOG TABLES.**
Reason: FDREPORT was unable to obtain the name of the master catalog.
Action: Contact Innovation for assistance.

100.03 CONTINUED . . .

- FDR688** **PROGRAM ERROR -- UNABLE TO error description**
Reason: The FDREPORT catalog processor failed for the reason displayed.
Action: Contact Innovation for assistance.
- FDR690** **pgmname ERROR CODE=nnnnnnnn**
Reason: FDRTSEL received a non-zero code upon return from the program whose name appears in the message.
Action: Examine the accompanying error messages generated by the failing program, contact Innovation if assistance is required.
- FDR691** **FDRTSEL SELECTING nnnnn UNIQUE TAPE VOLUME SERIALS.**
FDR691 **FDRTSEL VOLUME TABLE. THE FOLLOWING nnnnn VOLUMES**
WILL BE REQUESTED:
Reason: These are informational messages indicating the number files that FDRTSEL has selected for copying the number of unique tape volumes that will be allocated and the number tape mounts that will be performed based on the sequence of the tape mounts. A list of TAPE volume serial numbers, in the order they will be requested immediately follows these messages.
- FDR692** **type BACKUP SELECTED DSN=dsname VOL=volser**
Reason: The ARCHIVE or CATLG backup data set name shown in the message was selected for processing.
- FDR693** **ARCHIVE FILE RECORD UPDATED -**
VOL=volser ADATE=yyyy.ddd DSNAME=dsname
Reason: This message indicates that the archive file record containing the archive backup (that was copied or moved) was updated to reflect the new archive backup volumes. This message is output for all data sets contained in the archive backup.
- FDR694** **REASON=reason**
Reason: The checkpoint/restart function of FDRTSEL was requested but it failed. Reason codes are:
A - OPEN FOR DD TSELCKPT FAILED
Check the joblog of the FDRTSEL job for IBM messages. Check that the TSELCKPT DD is correct.
B - RESTART REQUESTED AND CHKPT FILE IS EMPTY
C - RE-OPEN FOR DD TSELCKPT FAILED
D - TSELCKPT FILE ID MISMATCH
An internal ID in the checkpoint file was incorrect. The file has been corrupted.
E - PREMATURE END OF FILE ENCOUNTERED ON CHECKPOINT FILE
The checkpoint file was not completely written or has been corrupted.
F - CONTROL STATEMENT VERIFICATION ERROR, SOURCE DOES NOT MATCH RESTART
Control statements used during the restart do not match those of the original FDRTSEL job.
G - RESTART RECORD ERROR, RECORD DOES NOT BEGIN WITH SELT, SELX, DSNR OR COMP
An internal error in the restart records was encountered. The file has been corrupted.
I - RESTART RECORD SEQUENCE ERROR code
An internal error in the restart records was encountered. "code" further describes the error.
H - OBTAIN FAILED FOR CHECK POINT DATA SET
The checkpoint file was not on the disk volume indicated by the JCL.
J - RESTART ARCHIVE CONTROL FILE OR VOLSER DOES NOT MATCH ORIGINAL
The name or volume serial of the Archive Control File used during the restart does not match that of the original FDRTSEL job.
Action: If possible, correct the error and resubmit. If necessary, contact Innovation for further assistance.
- FDR695** **ARCHIVE CONTOL FILE IS: ACfname**
Reason: This message documents the name of the Archive Control File used in this run.

100.04 MESSAGES FROM SAR (FDR9nn)

- FDR901 PROGRAM INTERRUPT -- JOB TERMINATED**
Reason: SAR detected a program check during its operation. A storage dump will be printed if a hardcopy device was assigned. The program check PSW is saved by SAR at location X'180'. This error may occur on a multi-processor CPU if the hardware START function was used to start any processor other than the one SAR was IPLed on.
Action: RelPL SAR. If you got a storage dump, and the error was not due to START, contact Innovation for assistance.
- FDR902 SVC INTERRUPT -- JOB TERMINATED**
Reason: SAR detected an SVC interruption during its operation; this should not occur. A storage dump will be printed if a hardcopy device was assigned. The SVC OLD PSW is saved by SAR at location X'180'.
Action: RelPL SAR. If you got a storage dump contact Innovation for assistance.
- FDR903 I/O ERROR -- JOB TERMINATED**
Reason: An unrecoverable I/O error occurred on the disk or tape. A FDR950-FDR990 message is printed detailing the error. A storage dump will be printed if a hardcopy device was assigned.
Action: If you can, try the operation again using different devices. If you cannot circumvent the error, contact Innovation for assistance.
- FDR905 TAPE DEBLOCKING ERROR -- REPLY CONTINUE OR TERM**
Reason: SAR detected an internal error in the tape block that was read. The tape block is at location X'6000' or X'20000' in storage.
Action: SAR will prompt you for a reply. Reply "CONTINUE" to ignore the error and continue with the next tape block, or "TERM" to terminate the restore with an error. If you have a hardcopy device assigned, you can use the hardware PSW RESTART function to take a storage dump before replying.
- FDR906 NOT AN FDR TAPE -- REPLY CONTINUE OR TERM**
Reason: The tape mounted for a SAR restore did not contain the FDR header. Location X'6000' in storage contains the bad header record.
Action: SAR will prompt you for a reply. Reply "CONTINUE" to ignore the error and continue with the next tape block, or "TERM" to terminate the restore with an error. If you have a hardcopy device assigned, you can use the hardware PSW RESTART function to take a storage dump before replying.
 1. If the tape is really an FDR tape which has been partially overlaid, replying "CONTINUE" will allow SAR to attempt to continue the restore without the header. Since the header contains the device type of the original disk, the user must ensure that the backup is being restored to the correct device type.
 2. Check that you have mounted the backup tapes in the proper sequence. If the tapes were not mounted in the proper sequence, reply "TERM" to terminate the restore and restart SAR.
- FDR907 TAPE BACKUP IS INCOMPATIBLE WITH DISK DEVICE TYPE -- JOB TERMINATED**
Reason: The backup tape was from a different device type than the restore device type (ex: SAR cannot restore a 3380 to a 3390). A storage dump will be printed if a hardcopy device was assigned. Location X'6000' in memory contains the FDR header.
Action: If you believe there was not a mismatch in the device types, contact Innovation for assistance in interpreting the header.
- FDR908 INVALID RESPONSE TO MESSAGE -- RE-ENTER WITH A VALID ADDRESS**
Reason: The response to the last message received was invalid. On a CRT device the cursor is repositioned. If response was for a device type, check for a supported device. If the response was for the unit field, check for valid hex characters (0-9,A-F). If WITH A VALID ADDRESS appears, check that the unit address is properly defined in your I/O configuration. If this is a response to the TAPE UNIT message, and you specified ",2" to the TAPE DEVICE message, check that the next higher tape drive address (or the next lower address if the address ended in "F") is also properly defined.
Action: Re-enter response to the message.
- FDR909 DISK VOLUME LABEL 'vvvvvv' DOES NOT MATCH OPERATOR LABEL -- RE-ENTER**
Reason: The volume serial number supplied by the user does not match the volume label on the disk. "vvvvvv" is the volume serial found on the DISK unit specified.
Action: Re-enter the proper volume serial number. Check the DISK unit address to insure you are pointing to the correct unit.

100.04 CONTINUED . . .

- FDR911 TAPE RECORD LENGTH CHECK -- REPLY CONTINUE OR TERM**
Reason: The length of the last tape block read does not match the length stored by FDR in the tape block. Location X'6000' or X'20000' in storage contains the tape block. If this message occurs multiple times, you may be restoring from an FDR backup which has been copied with a non-FDR utility such as IEBGENER; this will truncate the FDR data blocks and make the backup unusable. **Only the FDRTCOPY and FDRTSEL utilities (Section 60) or the Innovation FATAR product should be used to copy FDR tapes.** The input may also be a non-backup.
Action: SAR will prompt you for a reply. Reply "CONTINUE" to ignore the error and continue with the next tape block (the contents of the bad block will not be restored), or "TERM" to terminate the restore with an error. If you have a hardcopy device assigned, you can use the hardware PSW RESTART function to take a storage dump before replying.
- FDR913 WARNING BACKUP TAPE IS A DSF FORMAT TAPE -- REPLY CONTINUE OR TERM**
Reason: The FDR header on the backup tape mounted for a restore indicates that this was a data set backup, not a full-volume backup, and may not contain a copy of the entire volume. Location X'6000' in storage contains the FDR header.
Action: SAR will prompt you for a reply. Reply "CONTINUE" to ignore the error and restore the backup, or "TERM" to terminate the restore with an error. If you have a hardcopy device assigned, you can use the hardware PSW RESTART function to take a storage dump before replying.
- FDR914 NONSTANDARD RECORD ZERO -- X'cccchhhhrrkkllll'**
Reason: The record zero (R0) on the track indicated is of nonstandard size. The standard length is eight bytes. The count field of R0 is displayed: "cccc" and "hhhh" are the cylinder and track address of the bad R0. "rr" is the record number which should be zero. "kk" is the key length which should also be zero. "llll" is the data length which should be 0008; it is probably the data length which is wrong.
Action: If the record zero is shorter than 8 bytes, it will be dumped or restored as it appears on the disk. If it is longer, data in excess of 8 bytes will be lost.
- FDR916 LABEL ERROR ON TAPE REASON=x -- JOB TERMINATED**
Reason: SAR encountered an I/O error or logical error reading the tape or disk labels. May be accompanied by a FDR950-FDR990 message. If JOB TERMINATED does not appear, SAR will ask for a new tape to be mounted.
REASON CODE EXPLANATION
1 -- SAR did not find the HDR1 record on tape.
2 -- The file number encountered on the tape was higher than the file number requested.
3 -- Two tape marks were encountered trying to find the file number specified. The file is not on the tape.
4 -- While SAR was scanning the tape for the file number specified an end of volume (EOV1) record was encountered. The file number is not on the tape.
5 -- User specified a file sequence number other than zero or one and SAR found the tape to be non-labeled.
6 -- On a backup operation, the user mounted a non-labeled tape for a second or subsequent volume, when the first volume was labeled. Mount a SL tape.
7 -- An I/O error occurred reading the volume label on tape. Mount a new reel of tape.
8 -- CLIP was specified. SAR could not read the label on disk or the volume record was missing or incorrect. SAR cannot re-label the volume.
Action: Restart SAR and supply the proper file number if a standard labeled file. If file number is not known or not readable, use the non-labeled option (Nnn) of SAR to bypass the label and point directly to the data.
- FDR920 ENTER DATA SET NAME -- ENTER ***END TO PROCESS**
Reason: You specified the data set backup option. SAR is now requesting the data set names or groups to be dumped. SAR will display message DATA SET NAME= (on CRTs replacing the volume serial field). Enter the data set names after the equal sign. A data set name up to 44 characters may be entered.
If a group name is to be specified, end the group name with an *. SAR will dump all data sets found which match the characters specified before the *.
EXAMPLE: DATA SET NAME=SYS1* will dump all SYS1 data sets.
If the VTOC is to be dumped, enter ***VTOC. The will also dump cylinder zero head zero. A maximum of ten data set names or groups may be entered, one at a time.
Action: Enter the data set names to be dumped. Enter ***END to process the backup.

100.04 CONTINUED . . .

- FDR922 MOUNT NEXT REEL OF TAPE ON UNIT=uuuu**
Reason: The SAR operation requires that another reel of tape be mounted on the drive with address "uuuu".
 On Restore -- The backup file contains multiple reels of tape.
 On Dump -- The dump of the disk requires more than one reel of tape. Mount a scratch tape.
Action: Mount and ready the next reel. SAR will continue with the restore or backup operation, as soon as the tape becomes ready. If SAR does not automatically recognize that the tape has become ready, use the hardware PSW RESTART function to simulate the interrupt.
- FDR923 READY DEVICE nnnn UNIT=uuuu**
Reason: The tape, disk or hardcopy device with address "uuuu" indicated must be made READY before SAR can continue.
Action: Make the device READY. SAR will continue as soon as it becomes ready. If SAR does not automatically recognize that the device has become ready, use the hardware PSW RESTART function to simulate the interrupt.
- FDR925 SAR DISK VOLUME RELABELED TO vvvvvv**
Reason: You requested a clip operation. SAR has successfully relabeled the volume to "vvvvvv".
- FDR926 ENTER NEW VOLUME SERIAL NUMBER**
Reason: You requested a CLIP or new label (CPY=C). SAR is requesting the new volume serial after validating the volume serial now on the disk.
Action: Enter the new volume serial in the volume serial message.
- FDR927 ENTER ABSOLUTE TRACK ADDRESS IN DECIMAL -- ENTER ***END TO PROCESS**
Reason: You requested an absolute track operation. message STARTING CCHH=CCCCCHHHH will be displayed.
Action: Replace the CCCCCCHHHH with the cylinder (5 digits) and head/track (4 digits) in decimal. After ENTER is depressed SAR will display the ENDING CCHH= message; the ending cylinder and head/track is then entered. Up to 10 absolute ranges may be specified. When the last range is entered, reply ***END to the STARTING CCHH= message.
- FDR928 MAXIMUM NUMBER OF ENTRIES EXCEEDED -- ENTER ***END TO PROCESS**
Reason: You have entered the maximum number of entries (10) for data set names or absolute track ranges.
Action: You must enter ***END for SAR to process the data sets or tracks specified.
- FDR929 RESTORED CYL=cccc TRK -- 0....5....ABCDEF0....5....ABCDEF**
Reason: On absolute track operations, SAR will display each cylinder processed ("cccc" is the cylinder number in hex). An 'X' is printed in the appropriate column for each track restored within the cylinder.
- FDR930 TRACK TABLE ERROR**
Reason: A cylinder and head address was encountered on the backup tape which exceeded the track capacity of the disk volume.
Action: If a hardcopy device was assigned, use the hardware PSW RESTART function to take a storage dump. Contact Innovation for assistance.
- FDR931 VALID RESPONSE=responses**
Reason: This is an informative message to indicate the responses which are valid for the current console message. SAR may not display all of the possible responses. [Section 30](#) of the FDR Manual details all of the allowed responses to a message.
Action: Enter the appropriate response.
- FDR931 VALID RESPONSE=xxxxxx -- IDENTIFIED BY SAR**
Reason: The device type "xxxxxx" was determined by SAR to be the actual type of the disk device you input. It should not be overridden, except as discussed under "Alternate Device Support" in [Section 15.01](#).
- FDR932 I/O ERROR ON DISK TRACK X'ccccchhhh' -- REPLY CONTINUE OR TERM**
Reason: A logical or physical hardware error occurred on the disk track at cylinder "cccc" and track "hhhh" (in hex). A FDR951-FDR990 error message may have been displayed detailing the error. SAR did not backup or restore the indicated track properly.
Action: SAR will prompt you for a reply. Reply "CONTINUE" to ignore the error and continue with the next disk track (the contents of the bad track will not be dumped or may be improperly restored), or "TERM" to terminate SAR with an error. If you have a hardcopy device assigned, you can use the hardware PSW RESTART function to take a storage dump before replying.

100.04 CONTINUED . . .

- FDR933** **TAPE MOUNTED FOR *usage* IS status DSN=*dsname***
REPLY YES FOR SAR TO USE IT AS *type* -- NO TO REJECT -- EXPDT=*yyddd*
Reason: SAR displays the current data set name (last 17 characters from the tape label) for each tape mounted. "usage" may be HARDCPY, DUMP, or RESTORE.
 "status" indicates whether SAR detected labels on the tape and can be:
 IS VOL=vvvvvv for a labeled tape or
 NOT LABELED for an unlabeled tape.
 For an output tape, "type" will be SL or NL and the expiration date from the tape labels is displayed "yyddd". The operator must decide if this tape can be safely used for output.
 For an input tape, only the first line of the message is displayed and no reply is necessary; it simply documents the tape volumes used for input.
Action: Reply 'YES' for SAR to use the tape for output. If you reply 'NO', SAR will dismount the tape and request a new tape be mounted.
- FDR934** **REPLY WITH NEW TAPE DATA SET NAME**
Reason: SAR is requesting the data set name to be used on a labeled output tape. SAR will display the name found on the tape (last 17 characters). This message will be issued for the first output tape mounted.
Action: If you simply press ENTER, SAR will use the original data set name. You can override this name by specifying a new data set name. Up to 44 characters may be specified. If the NEWNAME is shorter than the original name, the remainder must be blanked out with the space bar (**the ERASE EOF key must not be used**). SAR will use the last 17 characters of the name you specify for all output tapes used in this backup.
- FDR935** **ERROR READING THE VTOC ON DISK REASON=*n* JOB TERMINATED**
Reason: SAR encountered an error reading the VTOC on disk during a backup operation.
 REASON CODE DESCRIPTION
 1 -- I/O error occurred reading the disk volume label. This may be due to an uninitialized volume.
 2 -- A logical or hardware error was detected in the VTOC. A FDR950-FDR990 message may have been displayed detailing the error.
 3 -- First record in the VTOC is not a Format 4.
 4 -- More than one Format 4 record encountered in the VTOC.
 5 -- An empty VTOC track was found.
 6 -- The VTOC track did not deblock. The records found were not a multiple of 140 bytes.
 7 -- A dummy VTOC on cylinder 0 track 0 was encountered. This usually indicates a volume initialized for VM.
Action: If JOB TERMINATED does not appear, SAR will continue the dump operation, dumping all of the tracks on the volume.
- FDR936** **DISK NOT SUPPORTED FOR DUMP -- JOB TERMINATED**
Reason: SAR cannot dump the volume specified. SAR uses the read multiple count-key-data CCW which is not supported on this control unit.
Action: Should not occur except on very old disk hardware.
- FDR937** **COMPRESSION TECHNIQUE NOT SUPPORTED**
Reason: SAR detected that the tape to be restored was compressed in a format that this level of SAR does not support. The tape was created by a higher release of FDR than the release of SAR.
Action: This tape must be restored by a release of SAR that is compatible with the release of FDR that created the tape; or the tape can be restored by a high-enough release of FDR under an operating system, or the tape can be copied to uncompressed format with a high-enough release of FDRTCOPY under an operating system.
- FDR938** **PRESS ENTER ON CONSOLE TO RESTART**
Reason: SAR has terminated and can be restarted without an IPL by pressing the ENTER key on the SAR console (EOB on printer/keyboard consoles).
- FDR939** **TAPE DSNNAME MISMATCH DSN=*dsname***
Reason: After mounting a volume of a multi-volume labeled restore tape, SAR found that the data set name in the header labels of the new tape does not match the data set name from the first input tape. The second line of message FDR933 will also be issued. You may reply "YES" to accept the tape and restore it or "NO" to reject the tape and mount the correct volume.
- FDR940** **ERROR READING LABEL ON DISK -- ENTER 6 BLANKS TO BYPASS**
Reason: An I/O error occurred trying to read the disk volume (VOL1) label. If the text "- ENTER 6 BLANKS TO BYPASS" appears, you may blank out the volume serial to cause SAR to ignore the error and continue. This can be useful if you are restoring to disk volumes which have not been initialized.

100.04 CONTINUED . . .

FDR951-FDR979	Reason:	The following messages indicate logical or physical I/O errors on the disk, tape, or hardcopy device. Each message contains these fields: "uuuu" -- device address of failing device cc -- CCW opcode of the failing CCW, in hex. "statrcnt" -- bytes 4-7 of the CSW (Channel Status Word) for this I/O ("stat"=status, "rcnt"=residual byte count) "ssssssssssssssssssss" -- sense information -- first ten bytes "cccchhhh" -- track address for disk only ("cccc"=cylinder number, "hhhh"=track number).
	Action:	Tape errors will cause immediate termination of SAR. Most disk errors can be bypassed (see message FDR932). Hardcopy errors will simply disable the hardcopy. You should try using a different device. If it still fails, call Innovation for assistance (24 hours a day). The text of each message is detailed below.
FDR951	COMMAND REJECT <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The device rejected the CCW command that SAR sent to it. 1. If it is a disk device, check that the correct disk device type specified. 2. If the command (cc) is a X'5E', the control unit does not support the read multiple CKD command. This disk cannot be dumped with SAR. It is followed by FDR936.
FDR952	HRDCOPY DISABLED <i>uuuu cc statrcnt sssssssssssssssss</i>	
	Reason:	SAR detected an I/O error writing to the hardcopy device (printer or tape). SAR will continue processing with the hardcopy device disabled, writing messages only to the console device..
	Action:	See Section 15 for information on supported console devices. Verify that the proper hardcopy device type was specified.
FDR953	BUS OUT CHECK <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The device reported a BUS OUT PARITY CHECK, which usually indicates a problem with the path to the device.
FDR954	EQUIPMENT CHECK <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The device reported an EQUIPMENT CHECK, which usually indicates a hardware problem with the device or a defect in the medium (e.g., tape).
FDR955	DATA CHECK <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The device reported a DATA CHECK. When reading, it indicates that a block that was successfully written to the device can not be read back. When writing, it indicates that the internal checking of the device found that the block was not written correct, and the internal control unit error recovery was not able to recover from the error.
FDR956	OVERRUN <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The device reported an OVERRUN, which indicates that the device could not accept data at the rate that the channel was presenting it, or vice versa. This usually indicates some error in your I/O configuration.
FDR957	ASSIGN ELSEWHERE <i>uuuu cc statrcnt sssssssssssssssss</i>	
	Reason:	A cartridge tape drive reported that it was currently assigned to another system. Another another CPU or LPAR has the device online.
	Action:	Specify another tape drive or vary offline this drive on all other systems.
FDR959	END OF CYLINDER <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	A disk device reported that SAR read past the end of a cylinder. Should not occur.
FDR960	SUBCHANNEL ERROR <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	After SAR started an I/O to a device, the subchannel for that device showed no I/O active. Should not occur.
FDR961	NOT AVAILABLE <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	SAR found that the device's unit address was not a valid address or not operational from this CPU.
	Action:	Check your I/O configuration for the proper unit address. If the device is accessible only from another CPU or LPAR, IPL SAR on that system.
FDR963	WRITE INHIBITED <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	During a restore, the disk device reported that is not enabled for write. Some older devices have a READ-ONLY switch. VM mini-disks and some modern disks can be configured to make certain disks read-only. Some control units will mark devices read-only after certain serious errors occur.
FDR964	INVALID TRACK FORMAT <i>uuuu cc statrcnt sssssssssssssssss cccchhhh</i>	
	Reason:	The data on the track to be dumped or the backup track being restored exceeds the track capacity of the disk device. For backup, this means that some application tried to write more data to that track than the track will hold.

100.04 CONTINUED . . .

- FDR965 FILE PROTECTED *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: During a backup or hardcopy to tape, the tape device reported that it was not enabled for write. For cartridge tapes, a thumbwheel is used to put tapes into read-only mode. On round (3420-type) tapes, a WRITE PROTECT RING is removed to make them read-only.
Action: if you really want to use this tape for output, turn the thumbwheel or insert the WRITE ring and restart SAR. Otherwise, choose another output tape volume and restart SAR.
- FDR966 DASD END-OF-FILE *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: While SAR was verifying the disk volume label, the disk device reported an END-OF-FILE (record with a data length of 0). The disk volume may not have a proper volume label (see message FDR940). Use six blanks in the volume serial to bypass the label check.
- FDR967 NO RECORD FOUND *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The disk device reported that a required disk record was not found. If this error occurs when the volume label is being verified, the disk volume may not have a proper volume label (see message FDR940). Use six blanks in the volume serial to bypass the label check.
- FDR968 INVALID ERROR *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The device reported an error which SAR cannot interpret.
Action: Make sure that you have specified the correct device type for this device.
- FDR969 WRONG ERROR *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The device reported an error status which should not occur.
- FDR970 CHAN CTRL ERROR *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported a CHANNEL CONTROL CHECK. This indicates a hardware problem with the channel, channel cables, control unit, or I/O configuration.
- FDR971 INTERFACE ERROR *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported an INTERFACE CONTROL CHECK. This indicates a hardware problem with the channel, channel cables, control unit, or I/O configuration.
- FDR972 CHAN DATA CHECK *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported an CHANNEL DATA CHECK. This indicates a hardware problem with the channel or channel cables.
- FDR973 NOISE RECORD *uuuu cc statrcnt sssssssssssssssss***
Reason: The tape device reported reading a noise record, which is usually a record smaller than the minimum the device can handle. This is similar to a data check (see message FDR965).
- FDR974 PROGRAM CHECK *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported an CHANNEL PROGRAM CHECK. This indicates a hardware problem or it may indicate an error in the CCWs used by SAR.
- FDR975 PROTECTION CHECK *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported an CHANNEL PROTECTION CHECK. This should not occur with SAR.
- FDR976 UNIT NOT CAPABLE *uuuu cc statrcnt sssssssssssssssss***
Reason: A 3480 cartridge containing data compacted by the 3480 IDRC feature was mounted for restore, but the 3480 tape drive chosen did not have the IDRC feature installed.
Action: If possible, choose a 3480 drive on another control unit which has the IDRC feature.
- FDR977 INCORRECT LENGTH *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: SAR requested that the device transfer data with a predetermined length, but the device reported that the length requested by SAR is different from the actual length of the data. This may indicate hardware errors on the device.
- FDR978 CHAINING CHECK *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The channel reported an CHANNEL CHAINING CHECK. This indicates that an invalid CCW chain was issued and should not occur in SAR.
- FDR979 SEEK CHECK ERROR *uuuu cc statrcnt sssssssssssssssss cccchhhh***
Reason: The disk device reported that the seek address (cylinder and head number of the track to be accessed) was invalid.
Action: Verify that the correct device type and model were specified for the disk device.

100.05 MESSAGES FROM ALL PROGRAMS (FDR99n)

- FDR990** **INSTALLED Innovation TRIAL WILL EXPIRE ON yyyy.ddd**
Reason: The expiration date of a trial version of FDR has been extended successfully. It will now expire on date "yyyy.ddd".
 Required JCL:
 //EXTEND EXEC PGM=FDREXTND, PARM=xxxx
 //STEPLIB DD DISP=SHR,DSN=fdr.loadlib
 //SYSLIB DD DISP=SHR,DSN=fdr.loadlib
 //@BINDNOT DD DUMMY
 //SYSDIAG DD SYSOUT=*
 The PARM= will be supplied by Innovation.
- FDR997** **FDR ABNORMALLY TERMINATED VOL=vvvvvv**
Reason: This FDR subtask has encountered an error from which it cannot continue. A user ABEND is being issued.
Action: A message detailing the error is printed.
- FDR998** **program status ERRORS VOL=vvvvvv COMPAKT BYPASSED**
Reason: "program" ran to completion on volume "vvvvvv" but there were error or warning messages. "status" can be COMPLETED or BYPASSED; the latter is used for a SPLIT or SNAP if the operation did not complete successfully. A user ABEND U0888 or a condition code of 8 or higher will be issued at the completion of the step. If COMPAKT BYPASSED is printed, FDR was requested to COMPAKT the volume after a successful backup; this will be BYPASSED since FDR detected one or more errors.
Action: Previous message(s) describe the error; see those messages for further details.
- FDR999** **program SUCCESSFULLY COMPLETED VOL=vvvvvv**
Reason: "program" ran to completion on volume "vvvvvv" without errors.

100.06 COMPAKTOR MESSAGES (CPKnnnx)

- CPK301I INNOVATION DATA PROCESSING -- COMPAKTOR VER v.r/mmt**
DATE yyyy.ddd TIME hh.mm.ss PAGE nnn
Reason: General page header message for COMPAKTOR. See message FDR101 for details
- CPK304I PARM -- *parm-field data***
Reason: Displays data input via the PARM= JCL operand on the EXEC statement.
- CPK305I CARD -- *control statement image***
Reason: Displays all control statements input from SYSIN.
- CPK306I ***END OF CONTROL STATEMENTS*****
Reason: Indicates the end of input control statements.
- CPK307I NO CPK CONTROL STATEMENTS WERE FOUND**
Reason: Neither the PARM= operand nor the SYSIN data set contained valid COMPAKTOR control statements.
- CPK316I COMPAKTOR RESTORED nnnnnn TRACKS**
Reason: "nnnnnn" is the actual number of tracks that COMPAKTOR wrote to the output volume. On Fast COMPAKtion or a COMPAKT-from-backup (DUMP=YES) to the original volume, this may be considerably less than the allocated tracks on the volume, since data sets which have not moved will not be restored.
- CPK319I SIZEKEEP PERFORMANCE OPTION - nnnnnn TRKS KEPT**
FOR DSN AT RELATIVE STARTING TRK ttttt
SIZEKEEP PERFORMANCE OPTION - nnnnnn TRKS KEPT TOTAL
Reason: The SIZEKEEP= option was specified (or the default values from the FDR Global Option Table were used). The first form of the message is issued for each data set which was made unmovable by SIZEKEEP processing; you can identify the data set by looking up relative track "ttttt" in the BEFORE or AFTER map. The second form of the message shows the total tracks which were made unmovable by SIZEKEEP. If SIZEKEEP is in effect but no CPK319I messages are printed, then SIZEKEEP was unable to meet the objective for the percent reduction in the free space areas without making all data sets movable.
- CPK320I COMPAKTOR OPTIONS IN EFFECT -- options**
Reason: Issued to inform you of run options in effect.
- CPK321I COMPAKTOR function STARTED**
COMPAKTOR RECOVERY STARTED DUE TO FAILED COMPAKTION
Reason: Issued just prior to start of disk volume modification for RESTORE, RELEASE, or RECOVERY. If this message is not issued, the disk volume was not altered in any way. After the message has been issued, the disk volume being processed is unusable until message CPK322I is issued.
The second form of the message is issued when a FASTCPK fails and COMPAKTOR enters automatic recovery. It can also occur when COMPAKTOR discovers that a previous Fast COMPAKTION job did not complete, and did not perform automatic recovery.
- CPK322I COMPAKTOR function ENDED.**
COMPAKTOR RECOVERY ENDED -
DATA SETS RETURNED TO ORIGINAL LOCATIONS
COMPAKTOR RECOVERY ENDED - REORGANIZATION COMPLETE
Reason: Indicates that COMPAKTOR has completed the operation **RESTORE, RELEASE, or RECOVERY**, and the volume should be usable. Some cleanup operations (such as mapping the COMPAKted volume and rebuilding the VTOC Index) may still remain to be done, so errors may still occur.
On a RESTORE from an FDR backup tape, if the text **SOME BACKUP VOLUMES WERE NOT REQUIRED** appears, this is simply an assurance that you need not be concerned if COMPAKTOR did not call for all the volumes of the FDR backup; tracks at the end of the volume did not need to be rewritten, probably because of %FOLD=nn or SIZEKEEP=.
The second form of the message is issued when COMPAKTOR RECOVERY has completed by backing out all changes and returning all data to its original location on the disk.
The third form of the message is issued when COMPAKTOR RECOVERY was able to complete the original COMPAKTOR request and has moved all data to its intended location.
- CPK403E EXCESSIVE NUMBER OF DATA SET AND/OR DSGROUP NAMES INPUTTED**
Reason: Over 255 different names were found on SELECT control statements.
Action: Rerun with fewer SELECT control statements.

100.06 CONTINUED . . .

- CPK404W** **operand=value IS INVALID. A DEFAULT VALUE MAY BE ASSUMED**
Reason: The value "value" is invalid for the operand "operand".
Action: Check for valid data or numeric data for the keyword, as described in [Section 40](#). If COMPAKTOR ABENDED, fix the invalid variable and rerun, otherwise check that the assumed default value is acceptable.
- CPK405E** **THE FROMDD= PARAMETER IS INVALID, OR ITS DD CARD IS MISSING, OR IT DOES NOT SPECIFY A type DEVICE, OR VOL= WAS SPECIFIED**
Reason: You specified FROMDD= but
 1) the value was invalid
 2) the DDname you specifies does not appear in the COMPAKTOR JCL.
 3) the device type allocated to the DDname is not **TAPE** or **DISK**
 4) VOL= was also coded (you cannot use both).
Action: Correct the error and rerun.
- CPK406E** **INVALID OR INCONSISTENT USE OF KEYWORD - operand=**
Reason: The operand "operand" is not valid for the form of the DSNNAME statement you used, or conflicts with a mutually exclusive keyword, which was also specified.
Action: Correct the error and rerun.
- CPK407E** **LAST 'SEQ' COMMAND SPECIFIED A NULL SET, OR THE 'ENDSEQ' COMMAND IS NOT PRECEDED BY A 'SEQ' COMMAND**
Reason: No SELECT control statements were found between a SEQ and an ENDSEQ statement, or an ENDSEQ control statement is misplaced.
Action: Correct error and rerun.
- CPK408E** **LAST READ CONTROL STATEMENT IS OUT-OF-SEQUENCE**
Reason: A major statement is preceded by one or more minor statements, or a SELECT DSG= statement precedes a SELECT DSN/SEQ/ENDSEQ statement.
Action: Correct the error and rerun.
- CPK409E** **A REQUIRED OPERAND IS MISSING**
Reason: An operand, either positional or keyword, required by the statement is missing.
Action: Correct the error and rerun.
- CPK410E** **EITHER 'ENDSEQ' COMMAND MISSING, OR POS=, DSG OR FILTER REQUESTED WITHIN A SEQUENCE SET. ASSUMING END OF SET**
Reason: 1) An ENDSEQ statement, delimiting a sequenced set, is missing.
 2) A SELECT statement with a POS= or SCRATCH= operand was encountered within the range of a SEQ statement. You cannot code the POS= operand for individual members of a sequenced set.
 3) A SELECT DSG= statement was encountered within the range of a SEQ statement.
Action: An ENQSEQ statement is assumed before the incorrect statement. If this doesn't produce the correct results, correct the error and rerun.
- CPK411E** **THE FOLLOWING NAME WAS SPECIFIED MORE THAN ONCE -- dsname**
Reason: The data set name "dsname" was found on two or more SELECT statements. You may specify a particular data set name or data set name group, only once.
Action: Correct the error and rerun.
- CPK412W** **THE 'RLSE=' OPERAND WILL NOT BE HONORED WHEN 'POS=KEEP' IS ALSO SPECIFIED**
Reason: Specifying POS=KEEP makes data sets unmovable, so their unused tracks cannot be freed
Action: None. This is only a warning message.
- CPK413E** **VTOC POSITIONING/SIZING IS INVALID WHEN 'VTOC=NOCHANGE' OR %FOLD IS ALSO SPECIFIED**
Reason: A SELECT ***VTOC statement was input to move and/or resize the VTOC, but VTOC=COMPAKT was not specified on the COMPAKT statement. It is also invalid if %FOLD= is specified.
Action: Correct the error and rerun.
- CPK425E** **UNABLE TO PROCESS VOL=vvvvvv -- DUMMY VTOC OR CMS VOLUME**
Reason: COMPAKTOR found that the VTOC on volume "vvvvvv" started and ended on cylinder 0 head 0. This is the form of a dummy VTOC created on volumes formatted for VM.
Action: The dummy VTOC does not contain any useful information, so COMPAKTOR will bypass this volume.

100.06 CONTINUED . . .

- CPK426E UNABLE TO OPEN DDNAME=ddddddd -- POSSIBLE MISSING DD CARD**
Reason: COMPAKTOR was unable to open for access the DD statement whose name is "ddddddd".
 This usually occurs when the DD statement is missing.
Action: Supply the appropriate DD statement and rerun.
- CPK427E THE DEVICE TO COMPAKT TO -- UNIT=uuu, VOLSER=vvvvvv IS NOT SUPPORTED BY COMPAKTOR**
Reason: An unsupported disk device was referenced. [See Section 2.10](#) for a list of supported device types.
Action: If necessary, contact Innovation for technical assistance.
- CPK428E THE DEVICE DUMPED TO TAPE ON UNIT=uuu, VOLSER=vvvvvv, IS NOT SUPPORTED BY COMPAKTOR**
Reason: The internal device type code for the disk type on the backup tape is not among those supported.
 The backup tape may not be a valid FDR backup.
Action: Contact Innovation for technical assistance.
- CPK429E INPUT DEVICE ON UNIT=uuu, VOLSER=vvvvvv IS NOT SUPPORTED BY TYPE=function**
Reason: An unsupported disk device was referenced for **COMPAKTOR** or **TYPE=FASTCPK**. See [Section 2.10](#) for a list of supported device types.
 If **TYPE=FASTCPK** is printed, the disk control unit is an old model which does not support the commands used by Fast Compaction.
Action: Contact Innovation for technical assistance.
- CPK430E THE DEVICE TO RESTORE TO - devtype1 reason THE INPUT DEVICE - devtype2**
Reason: When doing a COMPAKT-from-backup, and the output disk (type "devtype1") is not the same as the input disk ("devtype2") the disks are incompatible devices. "reason" may be:
 IS INCOMPATIBLE WITH - the disks are different device types (e.g., 3380 and 3390)
 IS SMALLER THAN - the disks are the same type but the output disk is smaller than the input.
 This is a warning; as long as all the data fits on the output disk and there are no unmovable data sets above the end of the output disk, the run will be successful.
Action: Retry with a valid device as the target device.
- CPK431E ERROR TRYING TO DETERMINE DEVICE TYPE**
Reason: A DEVTYPE SVC issued to verify the device type received a non-zero return code.
Action: Contact Innovation for technical assistance.
- CPK432E 'SIZE=nnnnn' IS TOO LARGE**
Reason: The specified new VTOC size nnnnn is excessively large.
Action: Correct the error and rerun.
- CPK433E INVALID CCHH POSITION SPECIFIED -- cccchhhh**
Reason: A POS= operand was specified with a track address but
 1) The cylinder address "cccc" is invalid. It exceeds the device capacity.
 2) The track address "hhhh" is invalid. It exceeds the maximum number of tracks per cylinder.
 3) The given value is all zeroes.
Action: Correct the invalid address and rerun.
- CPK451E INVALID CONTINUATION OF A QUOTED STRING**
Reason: A quoted string was improperly continued onto the following control statement. It does not start at column 16 of the continuation statement.
Action: Correct the error and rerun.
- CPK452E EXPECTED CONTINUATION CARD NOT RECEIVED**
Reason: Continuation was indicated on a control statement, but no further control statements followed.
Action: Correct the error and rerun.
- CPK453E CONTINUATION COLUMN IS NON-BLANK, BUT LAST PARAMETER DOES NOT END WITH A COMMA OR IN COLUMN 071**
Reason: See rules for continuing operand fields in the Introduction of this manual.
Action: Correct the error and rerun.
- CPK454E INTERNAL PROGRAM ERROR. CALL VENDOR FOR ASSISTANCE**
Reason: Program software error.
Action: Contact Innovation DATA PROCESSING for assistance.

100.06 CONTINUED . . .

- CPK455E UNBALANCED PARENTHESES. EXPECTING A PARENTHESES AT COLUMN *nnn***
Reason: A left (open) parenthesis does not have a matching right (close) parenthesis.
Action: Correct the error and rerun.
- CPK456E INVALID USE OF DELIMITER AT COLUMN *nnn***
Reason: A delimiter character (blank, parenthesis, apostrophe, comma or equal sign) is used in violation of the COMPAKTOR statement language syntax.
Action: Correct the error and rerun.
- CPK457E LAST READ CARD CONTAINED NO DATA**
Reason: A completely blank control statement was input.
Action: Correct the error and rerun.
- CPK458E UNIDENTIFIED COMMAND AT COLUMN *nnn***
Reason: The statement starting at column *nnn* is not a valid COMPAKTOR statement.
Action: Correct the error and rerun.
- CPK459E INVALID KEYWORD -- *kkkkkkk***
Reason: The keyword, *kkkkkkk*, is not a valid keyword parameter for this statement.
Action: Correct the error and rerun.
- CPK460E EXCESSIVE NUMBER OF OCCURRENCES OF COMMAND -- *cccccccc***
Reason: The statement "*cccccccc*" has appeared more than its maximum number of occurrences. Check statement description in [Section 40](#) of this manual for the maximum allowed occurrences.
Action: Correct the error and rerun.
- CPK461E UNIDENTIFIED POSITIONAL PARAMETER AT COLUMN *nnn***
Reason: A positional parameter was detected starting at column "*nnn*", but all positional parameters accepted by the statement have already been input.
Action: Correct the error and rerun.
- CPK462E UNBALANCED APOSTROPHES. EXPECTING AN APOSTROPHE AT COLUMN *nnn***
Reason: A quoted string does not begin and end with an apostrophe or an apostrophe is misplaced or a single apostrophe within a quoted string is not correctly represented by two consecutive apostrophes.
Action: Correct the error and rerun.
- CPK463E DUPLICATE KEYWORD DETECTED -- *operand***
Reason: The operand "*operand*" appears more than once within the operand field of a particular statement.
Action: Correct the error and rerun.
- CPK464E EXCESSIVE PARAMETER DATA LENGTH FOUND AT COLUMN *nnn***
Reason: The variable data for a positional or keyword parameter exceeds its maximum allowed length.
Action: Correct the error and rerun.
- CPK465E EXCESSIVE SUBPARAMETERS SPECIFIED STARTING AT COLUMN *nnn***
Reason: Too many positional suboperands, within parenthesis, were detected for a particular keyword or positional parameter.
Action: Correct the error and rerun.
- CPK470W *reason* DSN=*dsname***
Reason: The data set "*dsname*" was not moved. "*reason*" can be:
IS ACTIVE OR UNMOVABLE - the data set is either in use by another task or is considered unmovable by COMPAKTOR ([see Section 40](#) for a list of unmovable data sets).
IS UNCHANGABLE - the extents of the data set cannot be combined (usually issued for a VSAM component whose characteristics inhibit combination).
DATA SET WAS NOT MOVED - the data set was not selected for movement.
This message is printed only if the PRINT=ALL operand was specified.
- CPK501E I/O ERROR ON UNIT=*uuu*, VOLSER=*vvvvvv***
Reason: Issued when a permanent I/O error is detected on the direct-access device on unit "*uuu*", with a volume serial number of "*vvvvvv*".
Action: This message is always followed by a COMPAKTOR mini-dump. In this mini-dump you find the IOB, DCB, UCB, DEB, CCW CHAIN and the last home address read. Examine the IOB sense bytes to determine the reason for the I/O error. If you need assistance, please contact Innovation.

100.06 CONTINUED . . .

- CPK502E I/O ERROR -- *synadaf info***
Reason: While executing a COMPAKT-from-backup, a permanent I/O error was detected reading the FDR backup with BSAM. The IBM macro SYNADef was issued to get a formatted string indicating the nature of the I/O error; that string is printed as "synadef info".
Action: Examine the SYNADAF message to determine the cause of the error. For additional assistance, please contact Innovation.
- CPK503E UNEXPECTED END-OF-FILE ON INPUT TAPE**
Reason: While executing a COMPAKT-from-backup, the physical end of the backup was found before the special FDR trailer record was read. It is possible that the FDR backup tape volumes were mounted out of order or the FDR dump did not successfully complete.
Action: If message CPK321I has been printed and message CPK322I has not, the output disk volume is probably unusable. If the error can be corrected, or a good backup can be located, do an FDR restore.
 Contact Innovation if you need assistance recovering from this error.
- CPK504E INPUT TAPE IS NOT AN FDR DUMP TAPE**
Reason: While executing a COMPAKT-from-backup, either the input tape was not created by an FDR dump, or FDR encountered errors in the VTOC when dumping (message FDR125).
Action: Insure tape is an FDR volume dump tape (DSF created tapes are not accepted by COMPAKTOR). Insure that tapes were mounted in the order they were created.
- CPK505E TAPE BLOCKSIZE ERROR. FIRST 64 BYTES OF BLOCK READ FOLLOW**
Reason: While executing a COMPAKT-from-backup, the number of bytes read from the backup for a certain block does not match the internal blocksize or the block read is invalid. This may be due to tape drive hardware errors or tape media problems and can also be caused by copying an FDR dump tape with other than the FDR Tape Copy Utility FDRTCOPY (or FATAR).
Action: If possible, clean the tape and tape drive. Retry using a different tape drive or a different backup. If it still fails, contact Innovation for assistance.
- CPK506E AUTHORIZATION CHECK FAILED**
Reason: The volume being COMPAKted is defined to RACF with DASDVOL class protection, and the userid for this job does not have ALTER authority.
Action: Authorize the userid to the volume, or run under an authorized userid.
- CPK507E INPUT TAPE HAS VTOC OR LABEL ERROR**
Reason: While executing a COMPAKT-from-backup, COMPAKTOR detected that the FDR backup does not contain a valid VTOC or a valid disk volume label.
Action: Use FDR to dump the volume again, and rerun COMPAKTOR.
- CPK508E CPK DOES NOT SUPPORT DUMPCOMPRESSED TRACK (DCT) FORMAT TAPE. USE FDR RESTORE.**
Reason: While executing a COMPAKT-from-backup, COMPAKTOR detected that the FDR backup was created using DCT=YES (DUMPCOMPRESSEDTRACK=YES). COMPAKTOR cannot restore a DCT backup.
Action: 1) rerun the backup without DCT=YES and rerun COMPAKTOR or
 2) restore the backup with FDR and use FAST COMPAKTOR to do a COMPAKT-in-place.

100.06 CONTINUED . . .

CPK512I FORMAT 4 INDICATES VOLUME OCCUPIES nnnnn CYLINDERS

Reason: The Format 4 DSCB (the VTOC Descriptor record) indicates that the volume has "nnnnn" cylinders, which is different from what COMPAKTOR expects for this disk type. This disk may be a VM mini-disk or a disk which contains a non-standard number of cylinders.

This message is also issued if a disk volume was initialized with the wrong number of data cylinders or the wrong number of alternate tracks in the Format 4 DSCB. This might occur if:

- 1) a volume was moved from a disk device with alternate tracks to one with no alternate tracks, or vice versa (an FDR DUMP/RESTORE or COPY does not have this problem but other copy utilities may)
- 2) a volume was initialized with other than ICKDSF, such as a utility from your disk vendor.

This message may be accompanied by message CPK519E because the incorrect counts may make COMPAKTOR think that a data set near the end of the volume occupies non-existent tracks; if so, fix the problem with the counts as described below and rerun COMPAKTOR to see if the CPK519E goes away.

Action: COMPAKTOR will honor the value from the F4 DSCB if it is less than what COMPAKTOR expects. Otherwise COMPAKTOR will use its expected value.

If the data or alternate count was wrong, you can correct it:

- 1) an ICKDSF REFORMAT REFVTOC can be used to update the VTOC with the proper values without disturbing existing data. However the REFVTOC operand requires at least ICKDSF R16 plus a PTF. Check with IBM for the requirements and usage instructions for REFVTOC.
- 2) If you can't do REFVTOC, move all data sets from the volume and reinitialize it with an ICKDSF minimal INIT.

If you can't do either, contact Innovation for assistance.

CPK513E INVALID REQUEST TO ELIMINATE UNUSED TRACKS FOR DATA SET- dsname. DATA SET IS EITHER UNMOVABLE, UNCHANGEABLE OR IS NOT A SEQUENTIAL OR PARTITIONED DATA SET

Reason: A SELECT statement specified "dsname" and requested that unused tracks or extents be freed. However, the data set is ineligible for this action.

Action: Correct control statement in error and rerun.

CPK514W POSSIBLE FORMAT 5 ERRORS IN VTOC

Reason: 1) Either no Format 5 DSCBs exist or they are incorrect.
2) The 'contamination' bit is on in the Format 4 DSCB, or the volume is a DOS (VSE) volume.

Action: If COMPAKting, none required since COMPAKTOR will fix this problem. If mapping or simulating, either fix the VTOC yourself or COMPAKT the volume so that COMPAKTOR fix it for you.

CPK515E NON-STANDARD VTOC DETECTED. FORMAT 4 DSCB IS NOT THE FIRST RECORD OF FIRST VTOC TRACK

Reason: VTOC does not conform to IBM operating system standards. The first DSCB in the VTOC is not a Format 4 DSCB. This can also occur if the VTOC contains more than one Format 4 DSCB, which is also invalid. COMPAKTOR cannot process this volume.

Action: COMPAKTOR will bypass this volume, which may have to be initialized. Contact Innovation for assistance.

CPK516E DSCB AT CCHHR(HEX) -- cccchhhrrr -- IS PART OF A BROKEN F1-F2-F3 DSCB CHAIN

Reason: Either a Format 1 or Format 2 DSCB does not point to a Format 3 when it should; or a Format 2 or Format 3 DSCB has no F1/F2 DSCB pointing to it. "cccchhhrrr" is the disk address of the DSCB (cylinder, track, and record, in hex).

Action: List the VTOC using the IBM IEHLIST utility, e.g.,

```
//LIST EXEC PGM=IEHLIST
//SYSPRINT DD SYSOUT=*
//DISK1 DD UNIT=3390,VOL=SER=vvvvvvv,DISP=OLD
LISTVTOC VOL=3390=vvvvvvv,DUMP
```

In that listing, the disk address of each DSCB is printed on the right side of the third line for that DSCB. Locate the DSCB in error and attempt to fix the error condition.

If the DSCB identified by CPK516E is a Format 3, you can zap the DSCB to all binary zeros and then run a COMPAKTion to correct the VTOC. Sample JCL to do this zap is in member F3CLEAR in the FDR Installation Control Library (ICL).

WARNING: Message CPK516E may occur for conditions other than unchained F3 DSCBs, although that is the most common cause. DO NOT run the zap job until you have verified that all DSCBs identified in CPK516E messages are F3 DSCBs that are not themselves chained to another F3 DSCB (their last 5 bytes must be zero).

100.06 CONTINUED . . .

- CPK517E DATA SET NOT FOUND, OWNS NO EXTENTS, OR IS UNMOVABLE -- *dsname***
Reason: Data set "*dsname*", specified on a SELECT DSNNAME statement, was not found, was a model DSCB (no space allocated), or was unmovable. [Section 40](#) lists the types of data sets that COMPAKTOR considers to be unmovable.
Action: Correct the error and rerun.
- CPK518E DSCB AT CCHHR(HEX) -- *ccccchhhrr* -- INVALID OR UNSUPPORTED DSCB TYPE**
Reason: A DSCB was read and was not a Format 0, 1, 2, 3, 4, 5 or 7 DSCB. COMPAKTOR does not support Format 6 DSCBs, which describe split cylinder data sets, no longer supported on MVS.
Action: If Format 6 DSCBs exist, you cannot COMPAKT, MAP or SIMULATE this volume. You may want to list the VTOC as described under message CPK516I.
- CPK519E DSCB AT CCHHR(HEX) -- *ccccchhhrr* -- HAS AN INVALID EXTENT ADDRESS**
Reason: One or more extent descriptions in a DSCB are invalid. "*ccccchhhrr*" is the disk address of the DSCB (cylinder, track, and record, in hex). This can occur if:
 1) an extent descriptor is all zeros
 2) the starting address of the extent is less than the ending address
 3) for an SUL (standard user label) data set, the first extent is not exactly 1 track
 4) The starting or ending cylinder number is greater than the number of cylinders on the volume. This may occur if COMPAKTOR cannot properly identify the type or capacity of the disk.
Action: This is a serious error and should be attended to immediately. Follow procedures described in [Section 40](#) on VTOC ERRORS.
 This error can also occur if the volume has been initialized with an incorrect number of data cylinders or alternate tracks and the data set is at or near the end of the volume. If you also receive message CPK512I, correct the error as described under that message and rerun COMPAKTOR to see if the CPK519E message still occurs.
- CPK520W DSCB AT CCHHR(HEX) -- *ccccchhhrr* -- HAS AN INVALID EXTENT SEQUENCE NO**
Reason: Extents owned by a data set are numbered from 0 to 127, ascending. A violation of this rule was detected.
Action: See message CPK519E.
- CPK521x DSCB AT CCHHR(HEX) -- *ccccchhhrr* -- HAS INCORRECT EXTENT DESCRIPTOR FLAGS**
Reason: One or more extent descriptions in a DSCB have invalid flags. "*ccccchhhrr*" is the disk address of the DSCB (cylinder, track, and record, in hex). This can occur if:
 1) Flags indicate that the extent begins and ends on a cylinder boundary, when in fact it does not (message CPK521W).
 2) A flag indicating that the extent is a user label extent (LABEL=(,SUL)) was found on, but the extent is not the first extent (message CPK521E).
 3) Flags indicate extent is a shared (split) cylinder, not supported by MVS or COMPAKTOR (message CPK521E)
Action: For problem (1), COMPAKTOR will fix the problem as long as the data set is moved or has free space released. For problems (2) and (3), list the VTOC as described under message CPK516I. Locate the DSCB in error and determine which problem it is, and superzap the bad DSCB to fix the problem. Contact Innovation for assistance.

100.06 CONTINUED . . .

- CPK522E DSCB AT CCHHR(HEX) -- cccchhhrr -- HAS AN INVALID CHAIN POINTER OR COUNT FIELD**
Reason: Either the count field (hardware record ID) of the chain pointer (pointer to the next DSCB for this data set) was invalid. "cccchhhrr" is the disk address of the DSCB (cylinder, track, and record, in hex). This can occur if:
 1) The count field of the specified DSCB does not match its physical location.
 2) The chain pointer in a Format 1 or 2 DSCB is invalid (does not point to a record in the VTOC).
 3) The chain pointer in a Format 3 DSCB is not zero for data sets which cannot exceed 16 extents.
Action: Use FDRDSF PRINT (Section 20) to print the track on which the DSCB resides and determine which field is in error. If the count field is in error, you have a very serious problem and should contact IBM. If the chain pointer in a Format 1 or 2 DSCB is invalid, you can attempt to superzap the DSCB.
- CPK523E DSCB AT CCHHR(HEX) -- cccchhhrr -- HAS AN INVALID NUMBER OF EXTENTS**
Reason: The extent count in a Format 1 DSCB was invalid. "cccchhhrr" is the disk address of the DSCB (cylinder, track, and record, in hex). This can occur if the extent count is greater than 128 for ICF VSAM clusters and some SMS-managed data sets, or greater than 16 for other data sets.
Action: This is a serious problem. Determine the error type and attempt to fix the DSCB in error.
- CPK524E DSCB AT CCHHR(HEX) -- cccchhhrr -- DESCRIBES AN EXTENT WHICH OVERLAPS SOME OTHER EXTENT. OVERLAPPING EXTENT'S LOW AND HIGH CCHH(HEX) -- cccchhh-cccchhh**
CPK525E DSCB AT CCHHR(HEX) -- cccchhhrr -- IS THE OTHER DSCB INVOLVED IN OVERLAP. OVERLAPPED EXTENT'S LOW AND HIGH CCHH(HEX) -- cccchhh-cccchhh
Reason: COMPAKTOR has detected that two extent descriptor fields, in Format 1 or 3 DSCBs, claim ownership of the same track or tracks. This is a serious error. The data may be incorrect in either or both data sets. "cccchhhrr" is the disk address of the DSCB (cylinder, track, and record, in hex). "cccchhh" describes the overlapping extents (cylinder and track, in hex).
Action: Determine the contents of the overlapped area and eliminate or fix one of the DSCBs involved. If one of the DSCBs is a Format 3 that has no F1/F2 pointing to it, then that is the one to be eliminated. Use SUPERZAP to zero out this DSCB (see message CPK516E).
- CPK526E DSCB AT CCHHR(HEX) -- cccchhhrr -- IS A SECOND FORMAT 4 DSCB**
Reason: Only one Format 4 DSCB is allowed in a valid VTOC. The identified DSCB is a second Format 4 DSCB. "cccchhhrr" is the disk address of the DSCB (cylinder, track, and record, in hex).
Action: This is a serious problem. SUPERZAP may be used to zero the invalid DSCB. Sample JCL to do this SUPERZAP is in member F4CLEAR in the FDR Installation Control Library (ICL).
- CPK527I THIS DATA SET WILL NOT BE RESTORED -- dsname**
Reason: An informative message from COMPAKT-from-backup. Used to provide a list of all data sets (such as temporary data sets) scratched by COMPAKTOR.
- CPK529I INDEXED VTOC EXISTS ON VOLUME=vvvvvv**
Reason: An active VTOC Index (SYS1.VTOCIX.Vvolser) was detected on the volume. For SIM or MAP this is an informative message. On COMPAKT-from-backup, indicates that COMPAKTOR will convert the volume to OS format prior to COMPAKTion. At the end of a successful COMPAKTion, COMPAKTOR will reconvert back to an indexed VTOC by invoking ICKDSF. For a Fast COMPAKTion, the VTOCIX will be updated directly with the new free space information.
- CPK530E CPKWORK MISSING INDEXED VTOC DISABLED**
Reason: For a COMPAKT-from-backup, COMPAKTOR found that the volume contained an active indexed VTOC but the CPKWORK DD statement was missing. COMPAKTOR will convert the volume to OS VTOC format.
Action: You must include a CPKWORK DD statement if COMPAKTOR is to convert the volume back to indexed format. You can execute the BUILDIX function of ICKDSF independently to convert back to an indexed VTOC.

100.06 CONTINUED . . .

- CPK531E IXSYSPRT DYNAMIC ALLOCATION FAILED INDEXED VTOC DISABLED**
Reason: COMPAKTOR attempted to allocate DDNAME IXSYSPRT to DUMMY so that ICKDSF could be invoked to rebuild the indexed VTOC. However, the dynamic allocation failed.
Action: The IXSYSPRT DD statement could be included in the COMPAKTOR execution JCL to avoid the problem. For this volume, the user need simply execute the BUILDIX function of ICKDSF to rebuild the indexed VTOC since the COMPAKTion is completed.
- CPK532E ICKDSF ABNORMALLY TERMINATED WITH Cnnnn**
Reason: COMPAKTOR invoked ICKDSF to rebuild the indexed VTOC. However, ICKDSF ended with a non-zero return code.
Action: If a DD statement for IXSYSPRT to SYSOUT was included, examine the ICKDSF messages. If not, look up the ICKDSF return code "cccc" in ICKDSF documentation. If the error can be ignored, execute ICKDSF to rebuild the indexed VTOC.
- CPK533I INDEXED VTOC SUCCESSFULLY status**
Reason: There was an active Indexed VTOC (SYS1.VTOCIX.Vvvvvv) on the volume being COMPAKTed. "status" can be:
UPDATED - COMPAKTOR directly updated the free space maps in the Indexed VTOC.
REBUILT - ICKDSF was successfully used to rebuild the indexed VTOC on the volume being COMPAKTed.
- CPK534E INTERNAL ERROR CONVERTING INDEXED VTOC REASON=n COMP=nnnn**
Reason: An internal error in COMPAKTOR occurred when rebuilding the indexed VTOC. "nnnn" may be a reason code from a system function such as CVAf. The reason code "n" indicates various errors detected by COMPAKTOR.
Action: Contact Innovation for assistance.
- CPK535W CPK CANNOT REBUILD THE INDEXED VTOC -- USER MUST REBUILD IT USING ICKDSF**
Reason: The volume serial of the COMPAKTed volume duplicates that of another online volume, and COMPAKTOR has placed it offline.
Action: Once you can bring the volume online (perhaps after varying the original volume offline), execute the BUILDIX function of ICKDSF to rebuild the VTOC index.
- CPK536I VOLUME=vvvvvv IS A SYSTEM MANAGED VOLUME-- IN INITIAL STATUS**
Reason: The VTOC of the disk volume "vvvvvv" indicates that it is in System Managed Storage (SMS) format. If "IN INITIAL STATUS" appears, the volume is not yet completely converted to SMS.
- CPK537E HOWEVER THE UCB DOES NOT INDICATE SMS**
Reason: This will appear after message CPK536I to indicate that the UCB of the input disk does not also indicate that it is SMS-managed. This may occur if the SMS volume is being accessed from a system which does not have SMS installed or active or if the volume is not currently defined to SMS as a member of a Storage Group.
Action: None. COMPAKTOR will continue, unless message CPK538E is issued. However, if this system does have SMS active, it may indicate a problem with the SMS status of this volume.
- CPK538E WARNING -- OUTPUT VOLUME IS NOT SMS COMPATIBLE - VOLUME BYPASSED**
Reason: For a COMPAKT-from-backup, the System Managed Storage status indicated in the VTOC of the input volume (SMS-managed or not managed) does not match the SMS status indicated in the UCB of the output disk. You can SIMULATE the COMPAKTion, but any other function will get the text VOLUME BYPASSED and the COMPAKTion will be not be done.
Action: Ensure that the output disk has the same SMS status as the input disk volume or FDR backup. If you want to do the COMPAKTion in spite of this error (e.g., when moving or restoring SMS volumes on a non-SMS system), specify the operand SMSROT=NONE on the COMPAKT statement.
- CPK540W FORMAT n DSCB AT CCHHR(HEX) -- cccchhhrrr -- DEFINES AN INVALID FREE AREA IN EXTENT DESCRIPTOR nn**
Reason: A free extent descriptor in a Format 5 or Format 7 DSCB is invalid. "cccchhhrrr" is the disk address of the DSCB (cylinder, track, and record, in hex). "nn" identifies the descriptor within that DSCB (the nn'th extent descriptor). This can occur if the extent values are invalid (e.g., point to tracks beyond the end of the volume) or point to tracks which are currently allocated to a data set.
Action: If you are actually COMPAKTing the volume, COMPAKTOR will automatically create accurate Format 5/7 DSCBs.

100.06 CONTINUED . . .

- CPK541W FORMAT *n* DSCBs LEAVE *nnnnnn* TRACKS UNACCOUNTED FOR**
Reason: Some tracks are not allocated to any data set are not described by any Format 5 or Format 7 DSCB.
Action: If you are actually COMPAKTING the volume, COMPAKTOR will automatically create accurate Format 5/7 DSCBs.
- CPK542E UNABLE TO HONOR REQUEST TO ALTER VTOC POSITION OR SIZE. VSAM DATA SETS EXIST ON VOLUME**
Reason: A SELECT ***VTOC statement was input but non-ICF VSAM clusters with non-zero secondary allocation quantities exist on volume to be COMPAKTEd.
Action: Remove the SELECT statement; or, if you need to alter the size or location of the VTOC on this volume, you must first EXPORT all objects in VSAM data spaces with non-zero secondary allocation quantities, and DELETE the data spaces. After COMPAKTion, you can IMPORT the objects.
- CPK543W INVALID REQUEST TO CHANGE VTOC POSITION OR SIZE. VSAM DATA SETS EXIST ON VOLUME. SIMULATION WILL CONTINUE**
Reason: A SELECT ***VTOC statement was input but non-ICF VSAM clusters with non-zero secondary allocation quantities exist on source volume.
Action: COMPAKTOR continues the simulation and performs any requested VTOC position or size changes. However, this message is to warn you that your request to alter the VTOC's position or size will not be honored during a real COMPAKTion (message CPK542E).
- CPK552E THE FOLLOWING DATA SET IS INELIGIBLE FOR POSITIONING -- *dsname*. IT IS EITHER UNMOVABLE, HAS A STANDARD-USER LABEL OR GIVEN CCHH POSITION IS INVALID FOR THIS DATA SET**
Reason: The data set "*dsname*" was named on a SELECT statement with POS= specified or included in a sequence set with POS=. The positioning requested is not possible because:
 1) it is marked as unmovable.
 2) it has a user labels (LABEL=(,SUL))
 3) it must be on a cylinder boundary but the given position is not on such a boundary.
 4) it was positioned so that, while its first track was within the capacity of the volume, one or more tracks at the end of the data set were outside the volume track boundaries.
Action: Correct the error and rerun.
- CPK553E ONE OR MORE ABSOLUTE POSITIONED DATA SETS OVERLAP A DATA SET AT CCCC.HH(DEC) -- *cccc.hh***
Reason: One or more SELECT or SEQ statements requested absolute positioning for data sets (POS=) but it could not be honored because:
 1) A data set would overlap an unmovable or previously positioned data set.
 2) If the cylinder/head address is beyond the end of the volume, then a data set or a sequenced set of data sets was positioned so that its end was outside the boundaries of the volume.
 "*cccc.hh*" is the cylinder and track numbers (in decimal) of the data set which would be overlapped.
Action: Consult the extents map produced by COMPAKTOR and fix your control statements

100.06 CONTINUED . . .

- CPK554E UNABLE TO COMPAKT THIS VOLUME REASON=reason**
Reason: COMPAKTOR was unable to process the volume for "reason" which may be:
RECOVERY LOG FULL - the FASTCPK recovery log is full. More tracks were moved than can be recorded in the recovery log. Remove operands that cause many tracks to be moved (such as SIZEKEEP=0, OBJECT=MINEXTS, DATA=ALL and POS=) and rerun. Can occur if the volume has very few free space areas before COMPAKTion; you can add CPKFREEX=nn to avoid COMPAKTing such volumes. FASTCPK will go into recovery and return the volume to its previous state.
INSUFFICIENT FREE STORAGE TRACKS - VOLUME NOT CHANGED - there are too few unallocated tracks on the volume for FASTCPK to successfully operate.
NO FREE DSCB FOR RECOVERY LOG - VOLUME NOT CHANGED - there are no free Format 0 DSCBs in the VTOC. FASTCPK requires at least one free DSCB for its recovery log.
NO FREE TRACK FOR RECOVERY LOG - VOLUME NOT CHANGED - there is no track on the volume which will be in free space both before and after the FASTCPK. FASTCPK requires at least one such track for its recovery log.
ALGORITHMS FAILED - VOLUME NOT CHANGED - none of the internal repositioning algorithms used by COMPAKTOR was able to develop a new layout for the volume. This usually occurs when you use POS= to position data sets or have a large number of unmovable or active data sets and COMPAKTOR was unable to find a location for a large data set. Reduce the number of POS= operands and rerun.
INSUFFICIENT STORAGE - VOLUME NOT CHANGED - the region was insufficient for FASTCPK. Increase the region size.
I/O ERROR - VOLUME NOT CHANGED - an I/O error occurred on the output volume during initialization.
I/O ERROR WRITING RECOVERY LOG - VOLUME NOT CHANGED - an I/O error occurred formatting the FASTCPK recovery log.
NON-STANDARD DATA SET NAME - VOLUME NOT CHANGED - the VTOC of the volume contains a data set with a non-standard data set name, such as imbedded blanks or non-standard characters. FASTCPK cannot handle such data sets.
Action: A U0888 ABEND or non-zero return code will be issued at step end to call attention to the error; if you don't want the ABEND/code, specify UNABLE=IGNORE on the CPK statement. Except for RECOVERY LOG FULL, the volume has not been changed; if you can correct the error, rerun the job.
For RECOVERY LOG FULL, COMPAKTOR will reverse the tracks already moved, returning the volume to its original state. If the recovery fails, contact Innovation for assistance.
- CPK555E VTOC TOO SMALL TO HOLD ALL DSCBS. INCREASE VTOC SIZE AND RETRY**
Reason: On a COMPAKT-from-backup, you changed the size of the VTOC (SELECT ***VTOC,SIZE=nnn), but it is now too small for all the DSCBs which must be restored into it.
Action: Consult the summary map output by COMPAKTOR to determine the minimum VTOC size you require. Correct control statement in error and rerun. Unless you are sure that no new data sets will be allocated on this volume, you should leave a generous amount of free space in the VTOC.
- CPK556W COMPAKTOR UNABLE TO IMPROVE VOLUME BY function**
Reason: COMPAKTOR cannot improve the volume by COMPAKTion or RELEASE. Perhaps it was recently COMPAKTed. For TYPE=RLSE, no space could be released.
Action: A U0888 ABEND or non-zero return code may be issued at step end to call attention to the warning; if you don't want the ABEND/code, specify UNABLE=IGNORE on the CPK statement. Since no changes have been made to the volume, COMPAKTOR will not produce the 'after' map. If COMPAKTing, the volume will not be COMPAKTed, but a SIMULATION will be done instead. If you wish to force COMPAKTOR to COMPAKT the volume specify any positioning statement.
- CPK557E UNABLE TO ALLOCATE SEQUENCED SET CONTAINING DATA SET -- dsname. THE SEQUENCED SET IS PROBABLY TOO LARGE**
Reason: After allocating all unmovable data sets and the VTOC, COMPAKTOR tried to allocate space for the sequenced set containing the named data set; it was unable to do this, because it could not find a single contiguous area large enough to hold the entire sequenced set.
Action: Remove some data sets from the sequenced set and rerun.

100.06 CONTINUED . . .

- CPK559W** *function* **BYPASSED DUE TO VOLUME reason**
Reason: **SIMULATION, COMPAKTION or RELEASE** was bypassed on this volume because:
NOT EXCEEDING CONDITIONAL COMPAKTION OPTIONS - one or more of the operands CPKFREEX, CPKDSNMX, CPKMULTX, or MINRLSE were specified, but none of the specified values were exceeded on the volume
BEING EXCLUDED - the volume was named in an EXCLUDE entry in the COMPAKTOR unmovable table.
Action: The volume will not be COMPAKTed.
- CPK560W** **CPK DID NOT FIND REQUESTED VOL/STORGRP=name**
Reason: The VOL= or STORGRP= operands were specified, requesting dynamic allocation of the volume or SMS storage groups named. Either the storage group is not defined to SMS, or the volumes in the storage group or the volumes specified by VOL= were not mounted on any online disk device.
Action: COMPAKTOR will proceed to the next requested volume, if any. If the volume serial or storage group name was misspelled or inaccurate, correct it and rerun COMPAKTOR.
- CPK561E** **DISK DYNAMIC ALLOCATION ERROR VOL=vvvvvv UCB=uuu COMP=comp CODE=code**
Reason: A dynamic allocation of disk volume "vvvvvv" on device "uuu" failed with completion code "comp" and reason code "code". For OS/390, these codes are documented in the IBM manual Authorized Assembler Services Guide. They can also be found in Appendix A of the ISPF online HELP.
Action: COMPAKTOR will proceed to the next requested volume, if any.
- CPK562I** **COMPAKTOR PROCESSING NEXT -- options**
Reason: Processing of more than one volume was requested by the VOL= or STORGRP= operands. This message appears instead of CPK320I for the second and subsequent volumes.
- CPK563E** *operand* **INVALID REASON = reason**
Reason: The VOL= or STORGRP= operand was specified, but it is invalid for the reason indicated, which may include:
CANNOT BE SPECIFIED UNDER FDR/ABR - not valid when COMPAKTOR is invoked by FDR or FDRABR.
NOT SUPPORTED ON NON-MVS
ONLY ONE VOLUME CAN BE SPECIFIED ON CPK TYPE=COMPAKT
NO VOLUMES FOUND TO BE ONLINE - volumes specified are not mounted on any online disk
DISK1 DD STATEMENT SPECIFIED - cannot be used if a DISK1 DD statement is also provided.
Action: If possible, correct the error and re-execute COMPAKTOR.
- CPK564I** **COMPAKTOR WILL RELEASE nnnnnn TRACKS -- LEAVING mmmmmm TRACKS DSN=dsname**
COMPAKTOR WILL RELEASE nnnnnn TRACKS FROM VOL=vvvvvv
Reason: For TYPE=RLSE only, the first format lists the number of tracks "nnnnnn" to be released from the data set "dsname" (leaving "mmmmmm" tracks allocated to the data set). The second format indicates the total tracks to be released on volume "vvvvvv".
- CPK581E** **ANOMALOUS DATA ON TAPE WHOSE DD-NAME IS dddddddd**
Reason: During a COMPAKT-from-backup, this error may be caused by a number of internal consistency checks of the data read from a FDR backup.
Action: Contact Innovation for assistance.
- CPK582E** **TRACK ttttt WAS NOT RESTORED**
Reason: I/O errors forced COMPAKTOR to skip certain DASD tracks and/or tape blocks during the restore. Usually preceded by one or more CPK501 and/or CPK502 messages.
Action: "ttttt" is the track number, relative to zero, of the track not restored. To determine which data set owns this track, you must locate the track in the BEFORE COMPAKTion map. Either recreate the data set, or retry the COMPAKTion using different DASD or tape devices.

100.06 CONTINUED . . .

- CPK583E ERROR PROCESSING VVDS DATA SET CCHHR=cccc hhhh rr REASON=n**
Reason: COMPAKTOR had an error reading or updating the VVDS on the volume being COMPAKTeD. "cccc hhhh rr" is the cylinder/head/record (in decimal) of the VVDS block causing the error. "n" is a reason code:
 1 VVDS block is invalidly formatted
 2 VVDS block length is invalid
 3 I/O error reading the VVDS
 4 unable to acquire storage to save the VVDS
 A VVR record type is invalid
 B Type 23 cell is missing from VVR
 C Extent number is zero in type 23 cell
 D tracks/extent is not a multiple of CASIZE
 E Bytes/extent is not a multiple of CAs/extent
 F Extent exceeds volume size
 G HI-ALLOC-RBA invalid
 H Extent Count has increased
 J Extent Count in Type 60 cell is invalid
 K Extent Count in Q record invalid
 L Duplicate Type Q record (imbedded index)
Action: Numeric reason codes will cause the whole VVDS block to be bypassed; some clusters may not be updated correctly. Alphabetic codes cause the VVR to be bypassed; only the cluster in that VVR will be affected. You should run an IDCAMS DIAGNOSE on the VVDS when any of these errors occur to identify errors in the VVDS. If DIAGNOSE indicates that the VVR only exists in the VVDS and not in the VTOC, then run an IDCAMS DELETE VVR to remove the orphan VVR entry from the VVDS. If the message contains an alphabetic reason code and the component does exist on the volume, then run an IDCAMS REPRO to backup the cluster, next DELETE the cluster, then re-DEFINE the cluster and finally REPRO the data back into the cluster. An IDCAMS EXPORT with PERMANENT followed by an IMPORT can also accomplish the same results as the REPRO / DELETE / re-DEFINE / REPRO.
- CPK584E ERROR UPDATING VVR COMP=ssss CODE=uuuu DSN=component**
Reason: COMPAKTOR was updating the VVDS on another volume because a multi-volume cluster with a piece on this volume was changed, and an error occurred. The CPL (Catalog Parameter List) that failed is printed following the message. For COMP values less than 4000, "ssss" is the return code and "uuuu" is the reason code as defined under message IDC3009I in the appropriate IBM messages Manual.
 "ssss" values over 4000 are generated by COMPAKTOR:
 4089 -- HI-ALLOC-RBA has gone negative
 4090 -- HI-ALLOC-RBA is not a multiple of CISIZE
 4091 -- HI-ALLOC-RBA was not returned
 4092 -- AMDSBCAT address was not returned
 4093 -- Type 23 cell track count went negative on the volume being COMPAKTeD
 4094 -- Type 23 cell extent count went negative on the volume being COMPAKTeD
 4095 -- Type 60 cell extent count went negative
Action: This VVR is bypassed. This probably indicates an error in the VVR, such as duplicate or orphan VVRs. You should run an IDCAMS DIAGNOSE on each VVDS involved.

100.06 CONTINUED . . .

- CPK585E** **UNABLE TO MOVE CCHH=X'ccccchhh' reason**
Reason: FAST COMPAKTion was unable to move track "ccccchhh".
Action: This is an internal error. Contact Innovation.
- CPK586W** **VOL=vvvvvv IS IN RECOVERY DUE TO A FAILED COMPAKTion AT hh.mm ON yy.ddd**
Reason: A Fast COMPAKTion (CPK TYPE=FASTCPK) on volume "vvvvvv" started at the indicated time ("hh.mm") and date ("yy.ddd") but that COMPAKTion did not complete for some reason, such as a system crash. This message is subsequently produced by any COMPAKTOR function against the failed volume, including MAP. The volume is not in a usable status; many data sets will not be usable until COMPAKTOR recovery is completed.
Action: The Fast COMPAKTion must be restarted so that it can recover and make the volume usable again. If the COMPAKTOR step getting the CPK586W message is a Fast COMPAKTion, COMPAKTOR will automatically do this recovery, after asking permission from the operator via a FDRW80 message (specify RECOVERY=YES on the CPK statement if you want to bypass the WTOR and immediately proceed with the recovery). If this is a MAP or SIM, the step will proceed but will get a U0888 ABEND at the end. If it is a CPK TYPE=CPK (COMPAKT from backup), it will be terminated immediately. **In either case, you must submit a Fast COMPAKTion against this volume to make it usable by returning all data tracks to their original locations.**
- CPK587W** **VOL=vvvvvv RECOVERY BYPASSED DUE TO OPERATOR OR CONTROL STATEMENT REQUEST**
Reason: A Fast COMPAKTion (CPK TYPE=FASTCPK) run on disk volume "vvvvvv" (see message CPK586W) did not complete. This Fast COMPAKTion job attempted to recover from the failure but either the operator replied NO to the FDRW80 WTOR or RECOVERY=NO was specified on the CPK statement. Either of these indicate that recovery of the failed Fast COMPAKTion is not to be done.
Action: **The volume is not in a usable state. Many data sets may be unusable. You must submit another Fast COMPAKTion job and allow it to recover from the failure.**
- CPK595E** **module TERMINATED ABNORMALLY WITH COMP CODE Ssss Uuuuu**
Reason: A COMPAKTOR subtask abnormally terminated with either a system ABEND code "sss" (in HEX), or a user ABEND code "uuuu" (in decimal). "module" will be FDR for a FDR full-volume backup (DUMP=YES) or FDRRESTR for the actual modification of the volume.
Action: For SYSTEM ABENDS, refer to the appropriate IBM System Codes Manual. For user ABENDs, look it up in [Section 100.08](#). If the module is FDR, no changes were made to the volume, but if it is FDRRESTR and message CPK321I was printed without CPK322I the volume may be unusable. If possible, correct the error and rerun. Contact Innovation for assistance.

100.06 CONTINUED . . .

- CPK600In** **COMPAKTOR *status type* RUN -- VOL=vvvvvv**
Reason: Informs you of the status of the run. "type" is the type of run: **MAPPING, SIMULATION, RELEASE or COMPAKTION**. "status" will be **SUCCESSFULLY COMPLETED** or **COMPLETED WITH ERRORS**; in the latter case, some warning messages were issued but the COMPAKTion was able to complete. "n" is an internal code indicating which volume reorganization algorithm was used to position the data sets.
- CPK601I** **RESTORED DEVICE IS UNIT=uuu, VOLSER=vvvvvv, DEVTYPE=ddddddd**
Reason: Identifies the COMPAKTed volume if an output disk was modified.
- CPK602A** **SERIOUS ERROR(S) DETECTED. COMPAKTOR ABENDING**
Reason: One or more previous error messages were issued and COMPAKTOR terminates abnormally.
Action: See the Action for the other error messages.
- CPK603A** **COMPAKTOR RESTORE CANCELLED *reason***
Reason: The COMPAKTion of the current volume was bypassed because:
SYSTEM RESIDENCE VOLUME - you cannot COMPAKT the current SYSRES (IPL volume)
BY CONSOLE OPERATOR - the operator replied NO to the FDRW80 WTOR message.
- CPK604A** **COMPAKTOR OPERATION TERMINATED BY CANCEL -- REMAINING VOLUMES BYPASSED**
Reason: The operator attempted to cancel the COMPAKTOR job and it was intercepted by COMPAKTOR CANCEL protection. The operator then replied "S" to the FDRW99 message, causing COMPAKTOR to terminate after completing processing on the current volume. If more volumes were to be processed in the step, they are bypassed.
- CPK606E** **SNAPSHOT REQUEST ERROR - *reason***
Reason: FAST COMPAKTOR was executed with SNAPSHOT=YES, but an error occurred while using Snapshot on an IBM RVA or StorageTek SVA/Iceberg. "reason" details the error.
Action: Contact Innovation Technical Support.
- CPK607E** **SENSE ERROR - UNKNOWN DEVICE OR DEVICE NOT ECAM CAPABLE**
SENSE=<32 bytes of *sense data in hex*>
Reason: FAST COMPAKTOR was executed with SNAPSHOT=YES, but the current installed microcode does not support the ECAM commands required for Snapshot, probably because it was not an IBM RVA or StorageTek SVA/Iceberg disk subsystem. This message is just informational; it is not normally printed unless a debugging option is enabled.
Action: The volume will be COMPAKTed with normal READ and WRITE CCWs.
- CPK608E** **DWASET TABLE EXHAUSTED**
Reason: FAST COMPAKTOR was executed with SNAPSHOT=YES, but an internal table overflowed while using Snapshot on an IBM RVA or StorageTek SVA/Iceberg.
Action: Contact Innovation Technical Support.
- CPK609E** **EMC COPY REQUEST ERROR - MICROCODE LEVEL NOT HIGH ENOUGH**
Reason: FAST COMPAKTOR was executed with EMCCOPY=YES, but the microcode level of the EMC Symmetrix specified was not high enough to support the internal EMC COPY function.
Action: Contact EMC to upgrade your Symmetrix microcode or contact Innovation Technical Support.

100.07 FDRREORG MESSAGES (FDRRnn/FDRSnn)

FDRR00	statistics Reason:	Provides statistics about the current FDRREORG step, such as number of volumes processed, number of data sets processed, number of PDS tracks reclaimed, etc.
FDRR01	ALL PROCESSING INITIATED FOR VOLUME GROUP <i>gggggg</i>, VOLUMES PROCESSED=<i>nnnn</i>, VOLUMES EXCLUDED=<i>nnnn</i>, DUPLICATE VOLUMES BYPASSED=<i>nnnn</i> Reason:	A SELECT statement selected a SMS storage group with STORGRP=. This message shows the volumes in that group which were processed, excluded, or duplicates of volumes already in FDRREORG's volume list.
FDRR02	PROCESSING COMPLETED FOR VOLUME <i>vvvvvv</i> -- CODE=<i>cccc</i> statistics PROCESSING COMPLETED FOR DATA SET <i>dsname</i> -- CODE=<i>cccc</i> Reason:	"cccc" indicates the completion code for processing on a given disk volume or data set. Additional statistics are displayed.
FDRR03**	SUBTASK <i>nn</i> ABENDED PROCESSING VOLUME <i>vvvvvv</i> -- CODE=<i>cccc</i> SUBTASK <i>nn</i> ABENDED -- CODE=<i>cccc</i> SUBTASK <i>nn</i> ABENDED PROCESSING DATA SET <i>dsname</i> -- CODE=<i>cccc</i> Reason: Action:	An internal FDRREORG subtask has ABENDED with system or user ABEND code "cccc". For a system ABEND, determine if the ABEND was caused by a user correctable problem such as insufficient memory and rerun the job. It may be necessary to execute a RECOVER statement to complete any incomplete re-loads. For a user ABEND, please refer to Section 100.08 for an explanation of the ABENDs issued by FDRREORG. It may be necessary to contact Innovation for assistance.
FDRR04**	ALL SUBTASKS HAVE ABENDED -- PROCESSING TERMINATED Reason: Action:	All internal FDRREORG subtasks have ABENDED. Refer to the FDRR03 messages that were issued for each subtask ABEND.
FDRR05**	STOPTIME IS LESS THAN CURRENT TIME -- PROCESSING TERMINATED Reason: Action:	The STOPTIME parameter specifies a time that is less than the current. Provide an acceptable STOPTIME and rerun the job.
FDRR06**	STOPTIME PROVIDES LESS THAN 1 MINUTE OF RUNTIME -- PROCESSING TERMINATED Reason: Action:	The STOPTIME parameter must allow for at least one minute of runtime. Provide an acceptable STOPTIME and rerun the job.
FDRR07**	VALID SELECT STATEMENT NOT FOUND -- PROCESSING TERMINATED Reason: Action:	A REORG or SIMULATE statement was specified without a required SELECT statement. Provide at least one SELECT statement and rerun the job.
FDRR08**	ALLDSN REQUIRED IF A DSN LIST IS NOT PROVIDED -- PROCESSING TERMINATED Reason: Action:	The SELECT statement did not contain either a CATDSN or DSN operand. To process all data sets selected by the flagged SELECT statement, include ALLDSN to indicate that this is your true intention. Otherwise, add CATDSN or DSN.
FDRR09**	VOLUME LIST NOT PROVIDED -- PROCESSING TERMINATED Reason: Action:	The VOL= operand is required unless CATDSN is used to select data sets. Provide the VOL= operand or use CATDSN instead of DSN.
FDRR10*	VOLUME <i>vvvvvv</i> DOES NOT HAVE A VTOC -- VOLUME BYPASSED Reason: Action:	The volume does not have a usable VTOC. If this is not a VM volume, the indicated volume is probably damaged and needs to be restored.
FDRR11*	RUNTIME/STOPTIME REACHED -- NO NEW WORK WILL BE STARTED Reason: Action:	The limit specified by the user has been reached. FDRREORG will complete all active processing and process the next REORG or SIMULATE statement.
FDRR12*	OPERATOR STOP COMMAND ISSUED -- NO NEW WORK WILL BE STARTED Reason: Action:	The operator issued a STOP (P) command for the FDRREORG job. FDRREORG will complete all active reorganizations and terminate.
FDRR13**	VOLLIST BUILD ERROR -- PROCESSING TERMINATED Reason: Action:	Internal error. Contact Innovation for assistance.

100.07 CONTINUED . . .

- FDRR14** VOLLIST SCAN ERROR -- PROCESSING TERMINATED**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR15** ERROR BUILDING DSN MASKS -- RC=ccc**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR16** UNIT NAME cccccccc NOT DEFINED TO SYSTEM**
Reason: A unit name provided by operands such as BACKUPUNIT= is not defined in the current I/O configuration.
Action: Correct the unit name and rerun the job. If you have POOLDASD from Boole&Babbage and are using undefined unit names to control pooling, change the POOLDASD option to YES in the FDRREORG option table.
- FDRR17** EDTINFO FAILED -- RC=ccc, REASON=rrr**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR18** UCBLOOK FAILED -- RC=ccc, REASON=rrr**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR19** UNIT NAME cccccccc CONTAINS INVALID DEVICE TYPES**
Reason: A unit name provided by operands such as BACKUPUNIT= contains devices that can not be used by FDRREORG.
Action: Provide a unit name that references TAPE or DASD devices and rerun the job.
- FDRR20** UNIT NAME cccccccc CONTAINS A MIXTURE OF DEVICE CLASSES**
Reason: A unit name provided by operands such as BACKUPUNIT= contains a mixture of device types, such as TAPE and DASD.
Action: Provide a unit name that refers to TAPE or DISK devices, but not both, and rerun the job.
- FDRR21** UNIT VERIFICATION INTERFACE ERROR -- R15=ccc**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR22** UNABLE TO CREATE SMS MANAGED BACKUP DATA SETS -- SMS NOT ACTIVE**
Reason: SMS is not active and the user included operands that indicated that data sets which are dynamically allocated by FDRREORG are to be managed by SMS.
Action: Activate SMS or switch to unit name allocations.
- FDRR23** SMS SUBSYSTEM CALL ERROR -- R15=nnnnnnnn**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR24** SMS INTERFACE ERROR -- RETURN CODE=cccccccc, REASON CODE=rrrrrrrr**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR25** CLASS *classname* NOT DEFINED TO SMS**
Reason: The SMS class name provided by the user is not defined to SMS.
Action: Provide a valid class name and rerun the job.
- FDRR26** NOREORG LIST INVALID IN MODULE FDRNORG**
Reason: The permanent exclude list which is defined by the NOREORG statements contained in the module FDRNORG is not in a valid format.
Action: This module must be maintained by running FDRREOZO to add and delete the NOREORG statements.
- FDRR27** INVALID HEADER IN MODULE FDRNORG**
Reason: The permanent exclude list which is defined by the NOREORG statements contained in the module FDRNORG is not in a valid format.
Action: This module must be maintained by running FDRREOZO to add and delete the NOREORG statements.
- FDRR28** PARMLIST ERROR BUILDING NOREORG LIST**
Reason: Internal error processing the NOREORG statements in module FDRNORG.
Action: Contact Innovation for assistance.

100.07 CONTINUED . . .

- FDRR29** INVALID NOREORG LIST KEYWORDS IN MODULE FDRNORG**
Reason: The permanent exclude list which is defined by the NOREORG statements contained in the module FDRNORG is not in a valid format.
Action: This module must be maintained by running FDRREOZO to add and delete the NOREORG statements.
- FDRR30 function STARTED DSN=*dsname* BKUP=*bkdsname***
Reason: Indicates that the indicated FDRREORG function has started. This message will only be issued for PDS data sets if members are actually moved. If "function" is BLDAIX, the message will have BASE=basecluster in place of BKUP=.
- FDRR31 function STARTED DSN=*dsname* BKUP=*bkdsname***
Reason: Indicates the completion of the REORG function for a data set. See message FDRR30.
Action: None
- FDRR32 function PROCESSING TERMINATED**
Reason: The indicated FDRREORG function has been terminated due to errors.
Action: Refer to the preceeding error messages.
- FDRR33** function FAILED -- CODE=*ccccc*,DSN=*dsname***
Reason: The indicated FDRREORG function has failed due to a system or user ABEND.
Action: Consult IBM manuals to identify the proper action for a system ABEND. Refer to [Section 100.08](#) for user ABENDs issued by FDRREORG.
- FDRR34* NO DATA SETS FOUND IN CATALOG**
Reason: The CATDSN keyword was specified and no data sets matching the user specified data set names or filters were found.
Action: If there are additional SELECT statements they will be processed and the REORG or SIMULATE function will be performed provided the additional SELECT statements do not encounter the same error. If this is the only SELECT statement, processing will be terminated.
- FDRR35** ERROR ATTACHING function SUBTASK -- RC=*ccc***
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRR36 DATA SET QUALIFIES FOR REORGANIZATION -- DSN=*dsname***
Reason: The data set has met the requirements specified by the user for reorganization and will be processed.
- FDRR37** VSAM function FAILED -- R15=*ccc*,CODE=*rrr*,DSN=*dsname***
Reason: A VSAM request indicated by "function" failed. "ccc" is the return code in R15 and "rrr" is the reason code from the ACB. This may indicate that the cluster "dsname" is corrupted.
Action: Refer to the VSAM Macro Instruction Reference for a description of the record management return and reason codes. You may wish to run an IDCAMS EXAMINE on the cluster to see if it reports any errors. If necessary, contact Innovation for assistance.
- FDRR38* UNABLE TO PERFORM LAST TAPE PROCESSING**
Reason: Due to the error described by the preceding message, LASTAPE processing is not possible.
Action: Correct the error described by the preceding and rerun the job.
- FDRR40** VOLUME PROCESSING TERMINATED**
Reason: Due to the error described by the preceding message, volume processing is terminating.
Action: Correct the error described by the preceding and rerun the job.
- FDRR41** ENVIRONMENTAL ERRORS LIMIT REACHED**
Reason: The number of errors specified by the MAXENVERR keyword has been reached.
Action: FDRREORG will complete all active work and terminate.
- FDRR42** SYSTEM ABEND LIMIT REACHED**
Reason: The number of errors specified by the MAXSYSERR keyword has been reached.
Action: FDRREORG will complete all active work and terminate.
- FDRR43** UNABLE TO ESTABLISH ESTAE -- R15=*ccc***
Reason: FDRREORG was unable to establish its abnormal termination exit.
Action: Contact Innovation for assistance.
- FDRR44* ALL RETRY ATTEMPTS FAILED FOR DSN=*dsname***
Reason: A value other than NONE was specified for the DSNRETRY keyword and FDRREORG was unable to allocate the indicated data set before processing completed.
Action: Rerun the job when the indicated data set(s) are not in use by another job.

100.07 CONTINUED . . .

- FDRR45 LASTAPE NOT REQUIRED -- NEXT FILE FOR THIS TASK WILL HAVE FILESEQ=1**
Reason: Last tape processing was requested but the next backup data set will be the first data set on a new tape volume. Since a new volume will be used for the next execution, a LASTAPE catalog entry is not created.
Action: None
- FDRR46** ERROR CATALOGING LASTAPE DATA SET -- R15=ccc,DSN=dsname**
Reason: A catalog error occurred when FDRREORG attempted to record the last tape information in the catalog.
Action: Contact Innovation for assistance.
- FDRR47** catfunction ERROR -- RETURN CODE ccc,REASON IGG0CLxx rrr, DSN=dsname**
Reason: A catalog error occurred executing catalog function "catfunction" for "dsname". "ccc" is the catalog return code, "rrr" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Refer to IBM message IDC3009I. If you can't correct the error, contact Innovation for assistance.
- FDRR48** DATA SET BYPASSED**
Reason: Due to the error indicated by the preceding message, FDRREORG has bypassed the data set.
Action: Correct the error described by the preceding and rerun the job.
- FDRR49** READ JFCB FAILED -- DSN=dsname**
Reason: A RDJFCB macro has failed.
Action: Contact Innovation for assistance.
- FDRR50** VTOC OPEN FAILED**
Reason: An OPEN has failed for the VTOC.
Action: Refer to IBM IECnnn message in the job log.
- FDRR51** VVDS OPEN FAILED -- R15=ccc,CODE=rrr,DSN=dsname**
Reason: An OPEN has failed for the VVDS. "ccc" is the VSAM OPEN return code and "rrr" is the reason code from the ACB. This may indicate that the VVDS is damaged.
Action: Refer to the IBM manual *Macro Instructions for Data Sets* for a description of the VSAM OPEN return and reason codes. You may want to run a IDCAMS DIAGNOSE on the failing VVDS.
- FDRR52** OPEN FAILED FOR TARGET DATA SET -- R15=ccc,CODE=rrr, DSN=dsname**
Reason: An OPEN has failed for cluster "dsname". "ccc" is the VSAM OPEN return code and "rrr" is the reason code from the ACB. This may indicate that the cluster is damaged.
Action: Refer to the IBM manual *Macro Instructions for Data Sets* for a description of the VSAM OPEN return and reason codes.
- FDRR53** OPEN FAILED FOR BACKUP DATA SET -- DSN=dsname**
Reason: An OPEN has failed for the non-VSAM data set "dsname".
Action: Refer to IBM IECnnn message in the job log.
- FDRR54** CLOSE FAILED FOR TARGET DATA SET -- R15=ccc,CODE=rrr, DSN=dsname**
Reason: An CLOSE has failed for cluster "dsname". "ccc" is the VSAM OPEN return code and "rrr" is the reason code from the ACB.
Action: Refer to the IBM manual *Macro Instructions for Data Sets* for a description of the VSAM CLOSE return and reason codes.
- FDRR55** I/O ERROR READING VTOC -- ERROR CODE=ee,CSW=cccc,SENSE=ssss,CCCCHHHH=ccccchhh**
Reason: An I/O occurred while reading the VTOC. "ee" is the ECB post code, "cccc" is the channel status bytes from the CSW, "ssss" is the first 2 bytes of sense data from the disk device, and "ccccchhh" is the cylinder and track of the VTOC track on which the error occurred. All are in hex.
Action: Correct the VTOC or contact Innovation for assistance.
- FDRR56** VVDS GET FAILED -- R15=ccc,CODE=rrr**
Reason: An error occurred while reading the VVDS. "ccc" is the VSAM GET return code and "rrr" is the reason code from the ACB.
Action: Refer to the IBM manual *Macro Instructions for Data Sets* for a description of the VSAM GET return and reason codes.

100.07 CONTINUED . . .

- FDRR57** VVDS READ FAILED -- RC=ccc,RBA=xxxxxxx,VOL=vvvvvv,DSN=dsname**
Reason: A read request to the VVDS manager has failed. "cccc" is the catalog reason code from the VVDS manager and "xxxxxxx" is the RBA (Relative Byte Address) of the record being read.
Action: Refer to the IBM message manual for message IDC3009I, return code 50 and reason code "cccc".
- FDRR58** VVDS UPDATE FAILED -- RC=ccc,RBA=xxxxxxx,VOL=vvvvvv,DSN=dsname**
Reason: An update request to the VVDS manager has failed. "cccc" is the catalog reason code from the VVDS manager and "xxxxxxx" is the RBA (Relative Byte Address) of the record being updated.
Action: Refer to the IBM message manual for message IDC3009I, return code 50 and reason code "cccc".
- FDRR59** UNABLE TO ALLOCATE VVDS ON VOLUME vvvvvv**
Reason: Dynamic allocation of the VVDS has failed.
Action: Refer to the dynamic allocation failure message which follows this message.
- FDRR60** SYSDSN ENQ FAILURE -- RC=ccc,DSN=dsname**
Reason: An unexpected error occurred while attempting to enqueue the indicated data set for retry processing. "ccc" is the ENQ return code.
Action: Contact Innovation for assistance.
- FDRR61* DATA SET HAS MOVED FROM VOLUME vvvvvv TO VOLUME vvvvvv -- DSN=dsname**
Reason: A VSAM data set on the retry queue is not on the volume it was on when it was selected.
Action: The data set is removed from the retry queue and bypassed.
- FDRR62* DATA SET IS NO LONGER ON VOLUME vvvvvv -- DSN=dsname**
Reason: An IAM data set or PDS on the retry queue is no longer on the indicated volume. The data set has either been moved or deleted.
Action: The data set is removed from the retry queue and bypassed.
- FDRR63* RETRY DATA SET IS EMPTY -- DSN=dsname**
Reason: A data set on the retry queue has become available but is now empty. The indicated data set has probably been deleted and redefined.
Action: The data set is removed from the retry queue and bypassed.
- FDRR64* VOLUMES HAVE CHANGED FOR DSN=dsname**
Reason: An IAM data set or PDS on the retry queue has become available and the first volume of the data set has changed. The indicated data set has probably been deleted and redefined.
Action: The data set is removed from the retry queue and bypassed.
- FDRR65* DATA SET NO LONGER QUALIFIES FOR RE-ORGANIZATION -- DSN=dsname**
Reason: A data set on the retry queue has become available but no longer meets the selection criteria that caused it to be selected. The indicated data set has probably been reloaded or compressed by another job.
Action: The data set is removed from the retry queue and bypassed.
- FDRR66* DATA SET REMOVED FROM RETRY QUEUE**
Reason: Issued after a FDRR61, FDRR62, FDRR63, FDRR64, or FDRR65 message.
- FDRR67** FDRREORG WILL RETRY REORG LATER -- DSN=dsname**
Reason: A data set selected for reorganization is in use by another job and the SELECT statement that selected the data set specified a DSNRETRY option either RETRY, ENQ, or WAIT.
Action: The data set is added to the retry queue.
- FDRR68* MAXENQ LIMIT REACHED -- WILL ENQ LATER -- DSN=dsname**
Reason: A data set selected for reorganization is in use by another job and the SELECT statement that selected the data set specified a DSNRETRY option either ENQ or WAIT. An ENQ was not issued because the maximum number of outstanding ENQ's as specified by the MAXENQ keyword have already been issued.
Action: The data set is added to the retry queue as if the RETRY option was specified. An ENQ will be issued when the count of outstanding ENQ's falls below the value specified by MAXENQ.
- FDRR69** IAM 6.1 LEVEL 08 OR HIGHER IS EITHER INACTIVE OR NOT INSTALLED --PROCESSING OF IAM DATA SETS DISABLED**
Reason: The DSTYPE keyword specified IAM or ALL and a data set that appeared to be an IAM data set was encountered. An attempt to obtain information about the possible IAM data set using IAM's catalog interface indicated that IAM's catalog interface is not active.
Action: If IAM is installed, VIF must be active and it must be IAM 6.1 Level 08 or higher to reorganize IAM data sets. If IAM is not installed, do not use DSTYPE=ALL or DSTYPE=IAM.

100.07 CONTINUED . . .

- FDRR70** DATA SET HAS MORE THAN 16 EXTENTS -- DSN=*dsname***
Reason: The Format 1 DSCB for a partitioned data set has an extent count that exceeds 16.
Action: The Format 1 DSCB for the indicated data set is invalid and the data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR71** FORMAT 3 DSCB POINTER MISSING -- DSN=*dsname***
Reason: The Format 1 DSCB for a partitioned data set indicated that the data set is in more than 3 extents which requires that a Format 3 DSCB exist to describe the additional extents. The Format 3 pointer in the Format 1 DSCB however is zero.
Action: The Format 1 DSCB for the indicated data set is invalid and the data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR72** ERROR READING FORMAT 3 DSCB -- R15=*ccc*,DSN=*dsname***
Reason: An attempt was made to read the Format 3 DSCB for the indicated partitioned data set but it failed. "ccc" is the return code from the OBTAIN SVC.
Action: The Format 1 DSCB for the indicated data set is invalid or the VTOC has been damaged. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR73** FORMAT 3 DSCB NOT RETURNED BY OBTAIN -- DSN=*dsname***
Reason: An attempt was made to read the Format 3 DSCB for the indicated partitioned data set and the data return by OBTAIN was not a valid Format 3.
Action: Probable system error or VTOC damage. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR74** MISSING EXTENT DESCRIPTOR(S) -- DSN=*dsname***
Reason: A required extent descriptor for a partitioned data set in either a Format 1 or Format 3 DSCB was null.
Action: Probable system error or VTOC damage. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR75** EXTENTS OUT OF SEQUENCE -- DSN=*dsname***
Reason: While validating a partitioned data sets extent descriptors in the Format 1 or Format 3 DSCB, an extent sequence number was encountered that was not in the proper sequence.
Action: Probable system error or VTOC damage. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR76** INVALID FLAGS IN EXTENT DESCRIPTOR -- DSN=*dsname***
Reason: While validating a partitioned data sets extent descriptors in the Format 1 or Format 3 DSCBs, an invalid extent descriptor flag was encountered.
Action: Probable system error or VTOC damage. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors.
- FDRR77** INVALID CYLINDER/TRACK ADDRESS -- DSN=*dsname***
Reason: While validating a partitioned data sets extent descriptors in the Format 1 or Format 3 DSCBs, a cylinder or track address was encountered that is not valid for the device.
Action: Probable system error or VTOC damage. The data set will be bypassed. If you are licensed for COMPAKTOR, you can run a CPK MAP on the volume to identify all VTOC errors. See message CPK512E for a possible cause of this error.
- FDRR78** FORMAT 1 DSCB MISSING FOR DSN=*dsname***
Reason: The Format 1 DSCB for the indicated VSAM data component could not be found.
Action: Probable system error or VTOC damage. The data set will be bypassed.
- FDRR79** TRKCALC FAILURE -- RETURN CODE=*ccc*,DSN=*dsname***
Reason: A TRKCALC macro failed. "ccc" is the return code from TRKCALC.
Action: Contact Innovation for assistance.
- FDRR80** UNABLE TO FIND type ASSOCIATION FOR DSN=*dsname***
Reason: The association list for a VSAM cluster did not contain an entry for a required component type (DATA or INDEX).
Action: Probable catalog damage. Use an IDCAMS or TSO LISTCAT statement to verify the catalog entry for the indicated cluster. You may need to execute an IDCAMS DIAGNOSE on the catalog to identify the problem.

100.07 CONTINUED . . .

- FDRR81** ERROR SETTING INTEGRITY ATTRIBUTES -- R15=ccc, REASON IGG0CLxx rrr, DSN=dsname**
Reason: A catalog error occurred executing catalog function ALTER for "dsname". "ccc" is the catalog return code, "rrr" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Refer to IBM message IDC3009I. If you can't correct the error, contact Innovation for assistance.
- FDRR82** ERROR UNCATALOGING LASTAPE DSN -- R15=ccc,DSN=dsname**
Reason: An attempt to uncatalog the LASTAPE file "dsname" failed. "ccc" is the CAMLST catalog return code.
Action: Contact Innovation for assistance.
- FDRR83** ERROR CATALOGING BACKUP DATA SET -- R15=ccc, BACKUP DSN=dsname**
Reason: An attempt to catalog a tape backup data set failed. "ccc" is the CAMLST catalog return code.
Action: Contact Innovation for assistance.
- FDRR84** ERROR RESETING UPDATE INHIBIT FLAG -- R15=ccc, REASON IGG0CLxx rrr, DSN=dsname**
Reason: Before starting reload processing, an attempt to turn off the update inhibit flag turned on during backup processing has failed. "ccc" is the catalog return code, "rrr" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Contact Innovation for assistance.
- FDRR85** RE-USE OF DATA COMPONENT FAILED,DSN=dsname**
Reason: Unable to dynamically reuse a VSAM data component due to the error described by the preceding message.
Action: Refer to the preceding message and contact Innovation for assistance if necessary.
- FDRR86** RE-USE OF INDEX COMPONENT FAILED,DSN=dsname**
Reason: Unable to dynamically reuse a VSAM index component due to the error described by the preceding message.
Action: Refer to the preceding message and contact Innovation for assistance if necessary.
- FDRR87** ERROR UNCATALOGING BACKUP DATA SET -- R15=ccc, BACKUP DSN=dsname**
Reason: An attempt to uncatalog a tape backup data set failed. "ccc" is the CAMLST catalog return code.
Action: Contact Innovation for assistance.
- FDRR88** NON-GDGBASE CATALOG ENTRY EXISTS FOR DSN=dsname**
Reason: BACKUP=GDG was specified and the name generated by FDRREORG for the base generation data group is already in the catalog and is not a base generation data group.
Action: Remove the non-GDG entry from the catalog or use a different BACKUPGROUP or BACKUPINDEX.
- FDRR89** UNABLE TO DEFINE GDG-BASE -- R15=ccc, REASON IGG0CLxx rrr,DSN=dsname**
Reason: BACKUP=GDG was specified and an attempt to define the base generation data group has failed. "ccc" is the catalog return code, "rrr" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Contact Innovation for assistance.
- FDRR90** BACKUP DATA SET ALREADY CATALOGED -- DSN=dsname**
Reason: BACKUP=TEMP or BACKUP=PERM was specified and the backup data set name generated by FDRREORG is already cataloged.
Action: Delete the catalog entry or use a different BACKUPGROUP or BACKUPINDEX.
- FDRR91** CATALOG ENTRY TYPE FOR LAST TAPE DSN IS NOT NON-VSAM -- DSN=dsname**
Reason: LASTAPE was specified but the catalog entry for the LASTAPE name was marked as VSAM.
Action: After verifying that a data set name conflict does not exist between FDRREORG's LASTAPE data set name and a valid data set on your system, delete the existing catalog entry. Contact Innovation for assistance if necessary.

100.07 CONTINUED . . .

- FDRR92** ERROR RESTORING OWNERID AND ATTRIBUTES -- R15=ccc, REASON IGG0CLxx rrr DSN=dsname**
Reason: A catalog error occurred while FDRREORG was attempting to restore the original OWNERID and attributes of a VSAM data set. "ccc" is the catalog return code, "rrr" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Contact Innovation for assistance.
- FDRR93** ERROR EXTRACTING DYNALLOC MESSAGES -- R15=ccc**
Reason: A dynamic allocation error has occurred and an attempt to use the dynamic allocation error message interface has failed. "ccc" is the return code from the dynamic allocation message interface.
Action: Contact Innovation for assistance.
- FDRR94** UNABLE TO ALLOCATE type DATA SET- COMP=cc, CODE=rrrr, INFO=iiii, DSN=dsname**
Reason: A dynamic allocation error has occurred while attempting to allocate data set "dsname". "type" may be **TARGET**, **INDEX**, or **BACKUP**. "cc" is the dynamic allocation return code, "rrrr" is the reason code and "iiii" is the information code, all in hex.
Action: Refer to the IKJ message(s) that follows this message.
- FDRR96** UNABLE TO ALLOCATE VOLUME -- COMP=cc, CODE=rrrr, INFO=iiii, VOL=vvvvvv**
Reason: A dynamic allocation error has occurred while attempting to allocate volume "vvvvvv". "cc" is the dynamic allocation return code, "rrrr" is the reason code and "iiii" is the information code, all in hex.
Action: Refer to the IKJ message(s) that follows this message.
- FDRR97** DYNAMIC DEALLOCATION ERROR -- COMP=cc, CODE=rrrr, INFO=iiii, DDNAME=ddname**
Reason: A dynamic allocation error has occurred while attempting to deallocate "ddname". "cc" is the dynamic allocation return code, "rrrr" is the reason code and "iiii" is the information code, all in hex.
Action: Refer to the IKJ message(s) that follows this message.
- FDRR98** ENQ FAILED FOR IAM DATA SET -- DSN=dsname**
Reason: FDRREORG was unable to obtain the IAMENQ for the indicated data set because it is in use by another job.
Action: The data set is bypassed.
- FDRR99** ERROR GENERATING DSNAME FOR BACKUP DATA SET**
Reason: After applying the pattern as specified by the BACKUPINDEX keyword to the target data set name, an invalid data set name was detected.
Action: The data set name generated was probably longer than 44 characters. Simplify the BACKUPINDEX and rerun the job.
- FDRS01** ERROR OBTAINING NAME(+1) FOR GDG -- R15=ccc, DSN=dsname**
Reason: BACKUP=GDG was specified and a catalog error occurred when FDRREORG attempted to get the name of the +1 generation. "ccc" is the CAMLST LOCATE return code.
Action: Contact Innovation for assistance.
- FDRS02** DATA SET NAME TOO LONG TO BE A GDG-BASE -- DSN=dsname**
Reason: BACKUP=GDG was specified and the name of the generation data group is longer than 35 characters.
Action: Alter or use the BACKUPINDEX pattern to generate a shorter name.
- FDRS03** RE-READ OF JFCB FAILED FOR BACKUP DATA SET**
Reason: A RDJFCB macro has failed for the indicated backup data set after backup processing has completed.
Action: The target data set is not reloaded. Contact Innovation for assistance.
- FDRS04** ERROR LOCATING TIOT ENTRY FOR DDNAME=ddname**
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRS05** ERROR LOCATING JFCB IN SWA FOR DDNAME=ddname**
Reason: Internal error.
Action: Contact Innovation for assistance.

100.07 CONTINUED . . .

- FDRS06**** **UNSUPPORTED DEVICE FOR BACKUP DATA SET -- DEVICE CLASS=*cc*,DEVICE TYPE=*tt***
Reason: The BACKUPUNIT specified does refer to a disk or tape device.
Action: Specify a valid BACKUPUNIT and rerun the job.
- FDRS07**** **TARGET DATA SET NOT EMPTY -- DSN=*dsname***
Reason: A RECOVER statement was executed and the indicated data set has already been reloaded.
Action: The data set is bypassed.
- FDRS08**** **RE-USE TERMINATED**
Reason: An attempt to dynamically reuse a VSAM data or index component has been terminated due to the error described by the preceding message.
Action: Refer to the preceding message.
- FDRS09**** **REORGCKP DD REQUIRED FOR RECOVERY -- PROCESSING TERMINATED**
Reason: A RECOVER statement was executed without the JOBNAME keyword and a REORGCKP DD statement was not found.
Action: Provide the REORGCKP DD statement. If checkpoint processing is not required, use the NOCKPT keyword.
- FDRS10**** **REORGLOG DD REQUIRED FOR RECOVERY -- PROCESSING TERMINATED**
Reason: A RECOVER statement was executed without the JOBNAME keyword and a REORGLOG DD statement was not found.
Action: Provide the REORGLOG DD statement. If log file processing is not required, use the NOLOG keyword.
- FDRS11**** **RECOVERY TERMINATED**
Reason: A RECOVER statement has failed due to the errors described by the preceding message.
Action: Refer to the preceding message.
- FDRS12**** **DATA SET TYPE HAS CHANGED FROM OLD TYPE TO NEW TYPE**
Reason: A RECOVER statement was executed and the data set type has changed since FDRREORG backed up the data set.
Action: The data set is not recovered.
- FDRS13**** **TARGET DATA SET IS NOT A VSAM KSDS -- DSN=*dsname***
Reason: A RECOVER statement was executed and a target VSAM cluster is no longer a KSDS.
Action: The data set is not recovered.
- FDRS14**** **TARGET DATA SET IS NOT A IAM FILE -- DSN=*dsname***
Reason: A RECOVER statement was executed and an IAM data set is no longer in IAM format.
Action: The data set is not recovered.
- FDRS15** ***type* FILE ALLOCATED -- DSN=*dsname***
Reason: Documents the recovery checkpoint or log data set allocated.
- FDRS16**** ***type* FILE ALLOCATED ON UNSUPPORTED DEVICE**
Reason: The device allocated for the recovery checkpoint or log data set was not a DASD device.
Action: Specify a unit name that refers to DASD devices.
- FDRS17*** ***type* FILE DOES NOT EXIST FOR THIS JOB**
Reason: A RECOVER statement was executed with the JOBNAME operand but the required checkpoint or log data set does not exist.
Action: Probably none. FDRREORG will delete the checkpoint or log file if they do not contain an entry for a data set which requires recovery.
- FDRS18**** **RECOVERY NOT POSSIBLE -- A VALID CHECKPOINT OR LOG FILE COULD NOT BE FOUND**
Reason: A RECOVER statement was executed with the JOBNAME operand and FDRREORG could not find a checkpoint or log data set in the catalog.
Action: Verify that the specified jobname is correct.
- FDRS19*** ***type1* FILE IS OLDER THAN *type2* AND WILL NOT BE USED**
Reason: A RECOVER statement was executed with the JOBNAME keyword and FDRREORG found a checkpoint and a log file in the catalog. After inspecting the date and time stamps, FDRREORG determined that both files are not from the same job.
Action: FDRREORG will use the most current file only.

100.07 CONTINUED . . .

- FDRS20** RACF INFORMATION DOES NOT EXIST -- UNABLE TO SUBSTITUTE racfdata**
Reason: &RACFUID (USER ID) or &RACFGID (GROUP ID) was used with the BACKUPGROUP, BACKUPINDEX, CKPTPREFIX, LASTAPEPREFIX, or LOGPREFIX keywords and the user id or group id was missing from the ACEE control block.
Action: If you do not have a security system such as RACF installed, you may be able to utilize this feature if you are running MVS/ESA by providing a USER or GROUP parameter on the JOB card. If you do have a security system, then a system error has probably occurred. Have the system security administrator or system programmer check that the appropriate user and group information is being provided.
- FDRS21** type =UNIT NOT VALID -- operand IS NULL**
Reason: The operand **BACKUPALLOC=UNIT, CKPTALLOC=UNIT or LOGALLOC=UNIT** requested dynamic allocation of a data set, but the associated UNIT operand (**BACKUPUNIT=, CKPTUNIT= or LOGUNIT=**) was not specified and defaulted to NULL.
Action: Set a default value for the appropriate UNIT operand in the FDRREORG option table by running the FDRREOZO option table change utility, or specify a value at run time.
- FDRS22** type =SMS NOT VALID -- operand IS NULL**
Reason: The operand **BACKUPALLOC=SMS, CKPTALLOC=SMS or LOGALLOC=SMS** requested dynamic allocation of an SMS-managed data set, but the associated STORCLASS operand (**BACKUPSTORCLASS=, CKPTSTORCLASS= or LOGSTORCLASS=**) was not specified and defaulted to NULL.
Action: Set a default value for the appropriate STORCLASS keyword in the FDRREORG option table by running the FDRREOZO option table change utility, or specify a value at run time.
- FDRS23** BACKUPSTRING REQUIRES AN OLD AND NEW STRING -- SPECIFY BACKUPSTRING=(OLDSTR,NEWSTR)**
Reason: The BACKUPSTRING keyword was coded incorrectly.
Action: Correct the error and rerun the job.
- FDRS24** REORG OR SIMULATE COMMAND MUST PRECEED *statement***
Reason: A **SELECT** or **EXCLUDE** statement was encountered before a REORG or SIMULATE statement.
Action: SELECT and EXCLUDE statements must follow a REORG or SIMULATE statement. Correct the error and rerun the job.
- FDRS25** RECOVER COMMAND NOT SUPPORTED AFTER A REORG OR SIMULATE COMMAND**
Reason: Because of the basic differences in how the checkpoint and log files are used by a REORG and RECOVER function, it is not possible to perform both functions in a single execution.
Action: Run the RECOVER function as a separate job or jobstep.
- FDRS26** REORG OR SIMULATE COMMAND NOT SUPPORTED AFTER A RECOVER COMMAND**
Reason: Because of the basic differences in how the checkpoint and log files are used by a REORG and RECOVER function, it is not possible to perform both functions in a single execution.
Action: Run the REORG or SIMULATE function as a separate job or jobstep.
- FDRS27* PDS BYPASSED DUE TO UPDATEDPDS=NO -- DSN=*dsname***
Reason: The UPDATEDPDS option was specified or defaulted to NO and a partitioned data set was selected for compression that did not have a current backup (the update indicator in the data sets Format 1 DSCB was set).
Action: The data set is bypassed.
- FDRS28** type FILE NOT ALLOCATED ON A CYLINDER BOUNDARY**
Reason: The checkpoint and log files are dynamically allocated by FDRREORG as 1 cylinder data sets. However, after dynamically allocating either file, FDRREORG detected that they were not allocated on a cylinder boundary.
Action: Because of the additional overhead required to manage these files across cylinder boundaries, FDRREORG will terminate. The most probable cause for this error is a user allocation exit, or allocation control system, changing the allocation to a track allocation. You must allow the FDRREORG checkpoint and log file to be allocated as 1 cylinder data sets.

100.07 CONTINUED . . .

- FDRS29*** **CLUSTER OR AIX BYPASSED --
CAN NOT BE REUSED -- DSN=*dsname*
CLUSTER OR AIX CAN NOT BE REUSED
BUT WILL BE BACKED UP -- DSN=*dsname***
Reason: A VSAM KSDS or AIX was selected for reorganization and could not be reused. If a data set is not defined with the REUSE attribute, it can only be dynamically reused if it is a single volume data set with no related alternate indexes and not defined with keyranges. If VSAMDEFINE=IFREQ is specified, KSDSs that do not have RACF discrete profiles, VSAM passwords or keyranges will be deleted and re-defined by FDRREORG.
Action: If the ALWAYSBACKUP option was not specified, you will receive the first form of this message and the data set will be bypassed. If the ALWAYSBACKUP option was specified, the data set will be backed up but not reorganized.
- FDRS30*** ***function* NOT POSSIBLE -- DATA SET IN USE -- DSN=*dsname***
Reason: A partitioned data set selected for compression or simulated compression was in use by another job or user.
Action: The data set is bypassed. If PDSDISP was specified or defaulted to OLD, then any attempt at compression or simulated compression will fail if any allocations exist for the data set. If PDSDISP was specified as SHR, then some other user has it allocated exclusive with DISP=OLD.
- FDRS31*** **DATA SET HAS NEVER BEEN UPDATED -- BYPASSED -- DSN=*dsname***
Reason: A VSAM or IAM data set was selected for reorganization that had no adds, deletes, or updates, and the NOUPDATES option was specified or defaulted to NO.
Action: To avoid reorganizing data sets that do not require reorganization, the data set is bypassed.
- FDRS32*** **DATA SET HAS NOT BEEN LOADED -- BYPASSED -- DSN=*dsname***
Reason: A VSAM or IAM data set was selected for reorganization but the data set has been defined and never loaded.
Action: The data set is bypassed.
- FDRS33*** **UPDATE INHIBIT FLAG IS ON --
DATA SET BYPASSED -- DSN=*dsname*
UPDATE INHIBIT FLAG IS ON --
DATA SET WILL BE BACKED UP ONLY -- DSN=*dsname***
Reason: A VSAM data set was selected for reorganization and the data sets update inhibit flag is on.
Action: If the ALWAYSBACKUP option was not specified, you will receive the first form of this message and the data set will be bypassed. If the ALWAYSBACKUP option was specified, you will receive the second form of this message and the data set will be backed up but not reorganized. It is possible to receive this message if another FDRREORG is running and is processing this data set. FDRREORG sets the update inhibit flag to prevent the data set from being updated while it is being reorganized. If another FDRREORG is not running, then the update inhibit flag was probably set manually with an IDCAMS ALTER statement. The update inhibit flag will have to be reset with IDCAMS before FDRREORG will reorganize this data set.
- FDRS34*** **CLUSTER IS AN ESDS - WILL BE BACKED UP
BUT NOT REORGANIZED - DSN=*dsname***
Reason: The ALWAYSBACKUP option was specified or defaulted to YES and an ESDS (sequential) cluster was selected.
Action: The ESDS will be backed up.
- FDRS35**** **BACKUP DATA SET WOULD BE CATALOGED
IN THE MASTER CATALOG -- BACKUP DSN=*dsname***
Reason: The ALIASCHECK option was specified or defaulted to YES and a catalog alias does not exist to point the generated backup data set name to a user catalog.
Action: Backup processing is terminated.
- FDRS36**** **BACKUP=TEMP NOT ALLOWED WITH
ALWAYSBACKUP OR NOREORG**
Reason: The intent of the ALWAYSBACKUP or NOREORG options is to ensure that backups always exist for any data sets selected.
Action: Processing is terminated.

100.07 CONTINUED . . .

- FDRS37** RETRY ALREADY PENDING FOR THIS DSN ON ANOTHER VOLUME -- DSN=*dsname***
Reason: DSNRETRY was specified as ENQ or WAIT and a previous request for the same partitioned data set name had already been issued from another volume. Because the system does not manage SYSDSN ENQs on a volume basis, it is not possible to have more than one pending request.
Action: The data set will be bypassed.
- FDRS38** ERROR ALLOCATING VOLUMES FOR DEFINE -- DSN=*dsname***
Reason: An error occurred while allocating the volumes required to redefine a VSAM data set.
Action: Refer to the IKJ messages that follow this message and take the appropriate action. The RECOVER statement will have to be used to complete processing for the indicated data set. If there is a problem with the volume(s) required to successfully redefine the data set, the data set should be defined before using the RECOVER statement.
- FDRS39** VOLUME *vvvvvv* NOT MOUNTED -- DEFINE TERMINATED -- DSN=*dsname***
Reason: A volume required to redefine a VSAM data set is no longer online.
Action: Use the RECOVER statement to complete processing for the indicated data set when the required volume(s) are available. If there is a problem with the volume(s) required to successfully redefine the data set, the data set should be defined before using the RECOVER statement.
- FDRS40** UCBSCAN ERROR DURING DEFINE -- R15=*ccc*, REASON=*rrr*, DSN=*dsname***
Reason: A failure occurred using the UCBSCAN service. "*ccc*" is the return code from UCBSCAN and "*rrr*" is the reason code.
Action: Contact Innovation for assistance.
- FDRS41** RECORDS READ LESS THAN NUMBER OF RECORDS IN STATISTICS BLOCK -- DSN=*dsname***
RECORDS READ=*n*, RECORDS EXPECTED=*m*
Reason: To prevent reloading a damaged file, FDRREORG compares the number of records read during the backup ("*n*") with the number of records contained in the catalog statistics block ("*m*") for an IAM or VSAM data set. If the number of records read is less than the number of records in the statistics block, FDRREORG will not reload the data set.
Action: Research all activity against the indicated data set. If a system failure occurred while records were being deleted, it is possible that the statistics block was not updated to reflect the deleted records. If the data set is not damaged, FDRREORG can be forced to reorganize these data sets by specifying NORCOUNTERR on the REORG statement.
- FDRS43** UNEXPECTED DEVICE FOR BACKUP DATASET -- DEVICE CLASS=*cc*, DEVICE TYPE=*tt***
Reason: The POOLDASD option in the FDRREORG option table has been set to YES and an undefined unit name was used for allocation. The device allocated was not a DASD device.
Action: FDRREORG processing is terminated.
- FDRS44** UNABLE TO REDEFINE SMS MANAGED CLUSTER -- SMS NOT ACTIVE**
Reason: An attempt was made to redefine an SMS managed cluster and SMS is not active.
Action: Activate SMS and rerun the job.
- FDRS45* DATASET IS EMPTY AND WILL NOT BE REORGANIZED -- DSN=*dsname***
Reason: A previously loaded VSAM cluster or IAM data set no longer has any records.
Action: The data set is bypassed.
- FDRS46** NO DATASETS SELECTED**
Reason: No data sets matched the provided selection criteria.
Action: Check the SELECT/EXCLUDE statements and rerun the job.
- FDRS47** WAIT FOR DATA SET FAILED - ECB=*eeeeeeee*, DSN=*dsname***
Reason: DSNRETRY was specified or WAIT but the attempt to wait for the data set to become free failed. "*eeeeeeee*" is the ECB (Event Control Block) in hex.
Action: Contact Innovation for assistance.
- FDRS48** REVERTING TO ORIGINAL DEFINE -- DSN=*dsname***
Reason: An error occurred while attempting to redefine a VSAM cluster and operands were specified to change the original definition.
Action: FDRREORG will revert to the parameters used for the original define. Review and correct the requested changes and rerun the job to implement the desired changes.

100.07 CONTINUED . . .

- FDRS49**** **PRI=nnnnnn,SEC=nnnnnn,BLKSIZE=nnnnn, . . .**
Reason: Issued after message FDRR94 to list the parameters used for the failed dynamic allocation.
Action: Refer to the IKJ message(s) preceding this message.
- FDRS50**** **NONUNIQUE KEY DETECTED FOR UNIQUEKEY AIX -- DSN=dsname**
Reason: Multiple occurrences of the same alternate key has been detected for an Alternate Index defined with the UNIQUEKEY attribute.
Action: The alternate index build is terminated. This message is equivalent to the IDC1645I message issued by the BLDINDEX statement of IDCAMS.
- FDRS52**** **ALTERNATE INDEX HAS MORE THAN 3 PATHS -- DSN=dsname**
Reason: FDRREORG does not support Alternate Index data sets with more than 3 PATHs.
Action: The base cluster and all related objects are bypassed.
- FDRS53**** **EXCESS PRIME KEYS DETECTED DURING BLDAIX -- DSN=dsname**
Reason: The maximum recordsize of an Alternate Index was not large enough to contain all of the alternate keys.
Action: The alternate index is created with only the alternate keys that would fit. This message is equivalent to the IDC1646I message issued by the BLDINDEX statement of IDCAMS.
- FDRS54**** **AIX IN USE DSN=aixcluster BASE=basecluster**
Reason: A VSAM KSDS with an alternate index has been selected for reorganization and the alternate index is in use.
Action: The data set is bypassed.
- FDRS55** **statistics**
Reason: Informational message for IAM and VSAM data sets. This message is issued if RECORDCOUNTS=YES has been set in the FDRREORG option table.
- FDRS56**** **DUPLICATE KEY AT RECORD nnnnnnnnnn - DSN=dsname**
Reason: During backup of a KSDS or IAM data set, a duplicate key was found in record "nnnnnnnnnn" (decimal).
Action: The backup is terminated. You may need to reload the cluster or use other tools to fix the error.
- FDRS57**** **KEY SEQUENCE ERROR AT RECORD nnnnnnnnnn -- DSN=dsname**
Reason: During backup of a KSDS or IAM data set, an out-of-order key was found in record "nnnnnnnnnn" (decimal).
Action: The backup is terminated. You may need to reload the cluster or use other tools to fix the error.
- FDRS58*** **>4GIG KSDS CAN NOT BE REORGANIZED WITH THIS VERSION OF DF/SMS-- DSN=dsname**
Reason: During backup of a KSDS, FDRREORG detected that the cluster was defined as an "over 4 gigabyte" cluster on an SMS-managed cluster, but the level of the IBM DFSMS/MVS component installed does not support such clusters.
Action: The backup is terminated. Rerun on a system which includes DFSMS/MVS V1R3 or above (all OS/390 systems include V1R3 or above).
- FDRS59*** **SINGLE VOLUME DATA SET BYPASSED BECAUSE MODE=PARALLEL -- DSN=dsname**
Reason: Parallel backups were requested so single volume data sets are bypassed.
- FDRS60*** **MULTI VOLUME DATA SET BYPASSED BECAUSE MODE=SINGLE -- DSN=dsname**
Reason: Parallel backups were requested so single volume data sets are bypassed.
- FDRS61**** **type POINT FAILED -- R15=ccc,BACKUP=n,DSN=dsname**
Reason: During backup of a KSDS or IAM data set, a VSAM or IAM POINT failed with return code "ccc".
Action: The backup is terminated. You may need to reload the cluster or use other tools to fix the error.
- FDRS62**** **type READ FAILED -- R15=ccc,CODE=rrr,RBA=xxxxxxx,DSN=dsname**
Reason: During backup of a KSDS or IAM data set, a read for a DATA or INDEX record failed with return code "ccc" and reason code "rrr" on the record at RBA (Relative Byte Address) "xxxxxxx".
Action: The backup is terminated. Refer to the IBM manual Macro Instructions for Data Sets for a description of the VSAM return and reason codes. You may need to reload the cluster or use other tools to fix the error.

100.07 CONTINUED . . .

- FDRS63*** **CAN NOT DEFINE CLUSTER WITH IMBEDDED INDEX IN EXTENDED FORMAT -- DSN=*dsname***
Reason: The cluster has the IMBED option but the SMS data class assigned requests extended format (EF). IMBED cannot be used with extended format.
Action: The cluster will be allocated in non-extended format with the IMBED option. If you want to use extended format, specify the CONVERTINDEX operand and the cluster will be allocated in extended format with the NOIMBED option.
- FDRS64**** **BACKUP=GDG NOT SUPPORTED FOR PARALLEL BACKUPS**
Reason: MODE=PARALLEL was specified and GDG backups cannot be used.
Action: The backup is terminated.
- FDRS65**** **LASTAPE NOT SUPPORTED FOR PARALLEL BACKUPS**
Reason: MODE=PARALLEL was specified and the LASTAPE option cannot be used.
Action: The backup is terminated.
- FDRS66**** **? CHARACTER NOT FOUND IN BACKUP GROUP, INDEX, OR STRING - REQUIRED FOR PARALLEL BACKUPS**
Reason: MODE=PARALLEL was specified and BACKUPGROUP=, BACKUPINDEX= or BACKUPSTRING= operand must contain a ? to indicate where FDRREORG is to substitute the unique backup character.
Action: The backup is terminated.
- FDRS67**** **DEVICE TABLE ERROR - UNABLE TO CALCULATE ALLOCATED SPACE -- DSN=*dsname***
Reason: Internal error.
Action: Contact Innovation for assistance.
- FDRS68*** **DATASET TOO SMALL FOR MODE=PARALLEL - BYPASSED -- DSN=*dsname***
Reason: MODE=PARALLEL was specified but the data set is too small to be eligible for parallel processing.
Action: The data set is bypassed.
- FDRS69*** **IAM DATASET CAN NOT BE PROCESSED IN PARALLEL MODE - BYPASSED -- DSN=*dsname***
Reason: MODE=PARALLEL was specified but the data set is IAM.
Action: The data set is bypassed.
- FDRS70*** **IAM DATASET CAN NOT BE REDEFINED -- DSN=*dsname***
Reason: IAMDEFINE=YES was specified or default, but an IAM data set cannot be redefined. This may occur if you are running a level of IAM less than V6.3 or are reorganizing a data set that was created with a level of IAM less than V6.3. Other causes are possible.
Action: The data set is bypassed.
- FDRS71**** **VOLUME FLAGS NOT RETURNED -- DSN=*dsname***
Reason: A catalog LOCATE was issued for a cluster, but the volume flags returned by LOCATE were zero. This should not occur.
Action: Contact Innovation for assistance.
- FDRS72**** **OBTAIN FAILED -- R15=*ccc*,DSN=*dsname***
Reason: A DADSM OBTAIN was issued for an IAM data set, but it failed with return code "ccc".
Action: Contact Innovation for assistance.
- FDRS73** **REORG FORCED BY IAM -- DSN=*dsname***
Reason: IAM has recommended that an IAM data set be reorganized.
Action: Other conditional reorganization operands are ignored; the data set is reorganized.
- FDRS74**** **INCOMPLETE IAM INFO BLOCK - CAN NOT RECOVER FROM OUT OF SPACE -- DSN=*dsname***
Reason: An IAM data set does not have the information required to extend the data set.
Action: You may need to predefine the IAM data set and manually reload it from the backup created by FDRREORG.
- FDRS75**** **IAM HAS BEEN DEACTIVATED - CAN NOT DEFINE IAM FILE -- DSN=*dsname***
Reason: IAM is not active.
Action: Activate IAM and rerun the job.

100.07 CONTINUED . . .

- FDRS76** DUPLICATE VOLSER FOUND IN NEW type VOLUME LIST**
Reason: After redefining a VSAM cluster, duplicate volume serials were found in the volume list for the DATA or INDEX component.
Action: Contact Innovation for assistance.
- FDRS77** DUPLICATE type VOLUMES FOUND -- DSN=*dsname***
Reason: After redefining a VSAM cluster, duplicate volume serials were found in the volume list for the DATA or INDEX component.
Action: Contact Innovation for assistance.
- FDRS78* ERROR REMOVING DUMMY CANDIDATE VOLUME-
R15=*ccc*,REASON IGG0CLxx *rrr*,DSN=*dsname***
Reason: A catalog ALTER REMOVEVOL was issued but it failed. "*ccc*" is the catalog return code, "*rrr*" is the catalog reason code and IGG0CLxx is the IBM module reporting the error. The codes are documented under message IDC3009I in IBM message manuals.
Action: Contact Innovation for assistance.
- FDRS79** ADDVOL FAILED FOR IAM DATASET -- R15=*ccc*, DSN=*dsname***
Reason: An ADDVOL was issued for an IAM data set but it failed. "*ccc*" is the return code.
Action: Contact Innovation for assistance.
- FDRS80** REMOVEVOL FAILED FOR IAM DATASET -- R15=*ccc*, DSN=*dsname***
Reason: An REMOVEVOL was issued for an IAM data set but it failed. "*ccc*" is the return code.
Action: Contact Innovation for assistance.
- FDRS98 DATA SET HAS *nnnnnn* MEMBERS *nnnnnn*
DIRECTORY BLOCKS USED *nnnnnn* BLOCKS UNUSED**
Reason: FDRCOPY REORG was executed with the LIST=YES operand or REORG was invoked by the IEBCOPY interface and LIST=NO was not specified. The message gives PDS directory statistics.
- FDRS99 MEMBER *member status***
Reason: FDRCOPY REORG was executed with the LIST=YES operand or REORG was invoked by the IEBCOPY interface and LIST=NO was not specified. The message shows the status of each member in the PDS. "*status*" can be:
WAS MOVED - the member was relocated during compression
IN PLACE - the member did not need to be moved
IS ALIAS - member aliases are automatically updated when the prime member is moved.

100.08 ABEND CODES

Any of the programs and utilities in the FDR family may ABEND (abnormally end) with any of the following user ABEND codes. In many cases, a diagnostic message is printed before the ABEND, so look up any error messages that were printed first. If no message was printed that relates to this ABEND, then read the explanation below. Call Innovation if you need assistance understanding or resolving the error.

- U0100 Open Error Trying to Open a DASD or Tape DCB**
Usually preceded by a FDR324 message. Check the job log for IBM messages which may indicate the reason for the error.
- U0101 Maximum I/O Errors Exceeded on a Direct-Access Device**
More disk I/O errors than are allowed by the MAXERR= operand occurred. If you want to complete the operation in spite of the errors, specify a larger MAXERR= value. However, many data sets may be invalid.
- U0102 Alternate Tracks Invalidly Assigned**
Alternate tracks are not assigned according to IBM specifications. Can only occur during a dump without BUFNO=MAX.
- U0103 Disk Device Type Not Supported**
The type of disk being backed up or restore is not among those supported by FDR
- U0104 Full-Volume Restore to Unlike Disk Device Not Supported**
Use FDRDSF, FDRCOPY or FDRABR data set restore to restore to an unlike device.
- U0110 Invalid Extent Descriptor in DSCB or VVDS**
The descriptor for an extent of a data set in the DSCB or in the VVDS was invalid because it pointed to tracks outside the limits of the volume
- U0200 I/O Error on Backup Data Set**
A permanent I/O has occurred on the backup tape or disk data set. The backup is not usable. Can also occur if a RDJFCB, DEVTYPE, or OPEN fails.
- U0201 Premature End-of-File on Input**
The end of the backup file was encountered without finding the special FDR trailer record.
- U0202 Invalid Block read**
An invalidly formatted block was read from a backup.
- U0204 Invalid Block read**
An invalidly formatted block was read from a backup.
- U0205 Not an FDR Backup**
A restore detected that the backup was not created by FDR.
- U0207 Maximum Tape Block Length Checks Exceeded**
See message FDR204.
- U0210 Internal error in FDR Software Compression**
- U0300 Maximum for Count Field Errors on Direct-Access Device Exceeded**
More FDR123 messages than are allowed by the MAXERR= operand occurred. If you want to complete the operation in spite of the errors, specify a larger MAXERR= value. However, many data sets may be invalid.
- U0301 Maximum for Invalid Record Zeros on a Direct-Access Device Exceeded**
- U0302 Internal End-of-Extent Error**
The IGG019YZ module did not have the FDR identifier or was copied from a previous version of FDR or ABR. The most likely cause is that although you have installed a new version of FDR or ABR correctly in a Linklist library, there is an older version of IGG019YZ in LPALIB or in a library earlier in the Linklist concatenation. You must SCRATCH the older version.
- U0303 ABR Full-Volume Restore Cancelled, no VTOC on the Backup**
- U0304 Internal Error on Full Volume ABR Restore on Model DSCB**
- U0305 Invalid ABR Restore Allocation List (FDRALLOC)**
- U0401 SYSIN DD Statement Error**
SYSIN DD statement missing or incorrectly specified or I/O error on SYSIN data set.
- U0402 SYSPRINT/SYSPRINn/ABRMAP/SYSMAP DD Statement Error**
One of the above DD statements is missing or incorrectly coded or I/O error occurred processing the data set. There must always be SYSPRINT DD statement, and there must be a SYSPRINn DD statement for each backup TAPEn DD statement when using ATTACH or ABR.
- U0405 Invalid PARM field specified for FDR**
- U0502 One or more Control Statements are in Error**
- U0503 General Catalog Operation has Failed**
- U0504 Internal Sort Operations has Failed**

100.08 CONTINUED . . .

- U0600 Required DD Statement is Missing or in Error**
A message is always printed with the DDNAME and reason.
- U0601 Internal Error Reading a VTOC**
An FDR mini-dump will be printed.
- U0602 Error Restoring a Data Set or VTOC Error**
An FDR message will detail the error.
- U0603 Error on ABR Full Volume Restore**
message FDR328 was issued.
- U0604 Tape Volume Count Error**
More than nineteen (19) tape volumes were used for the backup file from one disk volume (for backup or archive), or POOLDISK wrapped around the pool and wrote over the beginning of the data set.
- U0605 VTOC Error Rewriting the ABR Model DSCB**
- U0608 FDRABR Tape Catalog Error**
The disk volume currently being dumped has been restored or modified by a program other than ABR (i.e.: an old release of FDR or SAR). An attempt to correct the condition and resynchronize the disk volume and the catalog failed.
- U0609 Dynamic Allocation Error**
message FDR336 shows the specific cause.
- U0610 Internal Wait Error**
This code can result from exceeding the maximum of nine unique units on TAPEX DD statements when running FDR or FDRDSF in ATTACH mode or FDRABR.
- U0611 Disk Device Type Not Supported**
- U0612 EXIT Error**
A parameter passed back by a user-written security exit was invalid.
- U0613 Error Determining the Device Type**
FDR issued a DEVTYPE SVC for a device and it failed.
- U0615 Last Tape Option Error**
ABR detected an ABEND condition trying to read the validation file for the last tape option. The IBM ABEND is printed on the job log. Resubmit the job and ABR will ask for a scratch tape.
- U0616 VTOC Error**
The number of DSCBs per track in the VTOC on a disk volume exceeded the maximum expected.
- U0619 ARCHIVE DSCB Error**
ABR attempted to ARCHIVE a data set that had invalid or overlapping extents.
- U0620 Internal Error on the Remote Queue Data Sets**
- U0621 Direct Access Device Type not Found in an Internal Table**
- U0623 ARCHIVE DSCB Error**
ABR attempted to ARCHIVE a data set that had an invalid FORMAT 2 or FORMAT 3 DSCB pointer.
- U0624 Internal DSCB Error on Format 2 or 3**
- U0625 TAPE or DISK DD Statement Error**
An ABR TAPEX or DISKxxxx does not point to a supported device type.
- U0626 An FDR/DSF/ABR Backup Format is in Error**
- U0628 Maximum backup files exceeded.**
See message FDR338.
- U0629 Internal Logic Error**
- U0632 Internal COMPAKTOR Unmovable Table is Full**
The table built by FDR when running under COMPAKTOR for unmovable data sets is full. The backup is terminated, and COMPAKTion is not attempted. The disk volume is not disturbed in any way. You can increase the size of the table (default: 8000) by specifying CPKUCORE=nnnnn on the DUMP TYPE=FDR statement. However, a large number of unmovable data sets will impact COMPAKTOR's ability to COMPAKT the volume, so the COMPAKTion may not be worth continuing.
- U0633 Internal Logic Error in BUFNO=MAX**
- U0634 More than one VVDS in the VTOC**
- U0650 DISKxxxx DD Statement OPEN Failed**
- U0654 VTOC Read Error**
- U0655 VTOC Rewrite Error**
- U0658 Mixed Program Levels**
See message FDR496.

100.08 CONTINUED . . .

- U0659 Internal Logic Error**
- U0701 Error in COMPAKTOR rebuild of indexed VTOC or update of VVDS**
- U0777 Internal program error, environment error or user error.**
Some possible causes are:
1. A module was called by a module from a different version or level.
2. A module received an invalid parameter list.
3. A TSO-only module was executed in batch.
4. An unexpected error was encountered in an Operating System routine.
Contact Innovation Technical Support for assistance.
- U0801 Restore Cancelled by Operator or security failure**
This ABEND can also be issued for other errors as described by messages in the listing. For an ABR SNAP or SPLIT with DSNENQ=, a data set was found to be in use by another task (message FDR 158) so the SNAP/SPLIT is cancelled; specify ENQERR=NO on the SNAP/SPLIT statement if you want to create the point-in-time backup anyway.
- U0802 Invalid Completion Code from a User Exit**
A parameter passed back by a user security exit was invalid.
- U0887 Message FDR321 REASON=E or M was issued.** An ABEND was forced by the DEBUG option.
- U0888 A Non-Terminating Error Occurred**
One or more FDR, DSF, FDRCOPY, ABR or COMPAKTOR operations abnormally terminated or ended with diagnostic messages, but the errors were not severe enough to prematurely terminate the entire step. This ABEND is issued to call attention to the errors. Examine the printout for the error messages causing this ABEND.
If a subtask abnormally terminated, the real ABEND code is printed in the storage dump and in the FDR319 message. FDR, DSF, FDRCOPY, and COMPAKTOR will issue this ABEND by default when non-terminating errors occur, but the ABEND can be changed to a return code by updating the FDR/ABR Global Option Table. ABR issues a return code of 12 in this case, by default, but it can be changed in the option table to issue the U0888 ABEND instead.
- U0901 Error executing the FDR trial extension program, FDREXTND.**
- U0902 Error executing the FDR trial extension program, FDREXTND.**
- U0996 Internal Debugging ABEND**
Contact Innovation if this ABEND occurs unexpectedly.
- U0997 Internal Debugging ABEND**
Contact Innovation if this ABEND occurs unexpectedly.
- U0999 Reload Failure in FDRREORG**

100.09 WAIT STATE CODES FROM SAR

SAR issues wait states indicating an action is required or an error has occurred. Most wait states are preceded by an explanatory message. The wait state code is the last four digits of the current PSW. On many CPUs the PSW is automatically displayed on the hardware console when the STOP key is pressed; consult your hardware operating manuals for details.

WAIT STATE**00E2 MACHINE CHECK**

Reason: A machine-check (CPU error) has occurred and SAR is unable to continue.

Action: Display storage location 232 (X'E8') to get the machine check interrupt code, which can be interpreted using information in the PRINCIPLES OF OPERATION and/or REFERENCE SUMMARY for your CPU. Contact your hardware maintenance vendor for assistance.

1111 I/O WAIT AND RETRY PSW

Reason: A normal wait state during I/O operations.

3333 CONSOLE ERROR

Reason: An I/O error has occurred reading or writing from the master console. This is not preceded by a message.

Action: RelPL SAR and retry using a different console.

5555 PROGRAM CHECK

Reason: A program check occurred before SAR identified the console. SAR was not able to issue a message.

Action: Contact Innovation for assistance.

9997 TERMINATE WAIT STATE

Reason: An I/O error or other error has abnormally terminated SAR. An FDR9xx message detailing the error should have been printed.

9999 EOJ WAIT STATE

Reason: Normal wait state indicating end of job. an FDR99x message should have been printed.

FFFF OPERATOR CONSOLE WAIT STATE

Reason: This wait state indicates that SAR has been successfully loaded. SAR is now waiting for the interrupt which will identify the device to be used as the console

THIS PAGE WAS INTENTIONALLY LEFT BLANK

**READER'S COMMENT FORM
FDR USERS MANUAL V5.3
INNOVATION DATA PROCESSING**

If you have comments on this manual, including:

- errors in the text or typographical errors
- clarity
- suggestions for improvement in the manual
- suggestions for improvement in the product
- any other comments

Please complete this form and fax it to Innovation at 973-890-7147 (in Europe you may fax it to your local Innovation office as shown on the front page of the manual).

You may also e-mail your comments to Innovation at support@fdrinnovation.com (be sure to identify the manual name in the message).

Your name: _____

Company name: _____

Mailing address: _____

E-mail address: _____

Comments: _____

THIS PAGE WAS INTENTIONALLY LEFT BLANK